



Center for
AI Oversight

The AI Oversight Baseline Review

A fourteen-domain self-review of the oversight layer of AI governance

FOR BOARDS, EXECUTIVE MANAGEMENT, AND THE FUNCTIONS THAT REPORT TO THEM

A FREE RESOURCE OF THE CENTER

What this review is, and what it is not

A baseline, not a maturity score

Most organizations that ask whether they govern AI are asking the wrong layer of the question. They can describe their policies and their model controls. They cannot say, in one sentence, who is accountable for AI oversight and what the board would show a regulator if asked.

This review asks that question fourteen times, once for each domain of the AI Oversight Program. For each domain it asks whether the oversight element exists, who is accountable for it, and what evidence would show it. It takes an afternoon. It requires no technical knowledge and no consultant.

It produces a baseline: a picture of whether an oversight program exists, where accountability is named, and where the organization does not know. It does not produce a maturity score. Maturity is a question of how deeply oversight has been institutionalized, and answering it takes evidence and a method. That is the work of the Center's AI Governance Maturity Model. This review tells you whether there is a program to assess.

Who should complete it

The executive accountable for AI oversight, with the general counsel, the chief risk officer, and the chief audit executive. A director may complete it alone to test what the board actually knows.

What it covers

The oversight layer only: mandate, accountability, appetite, standards, escalation, and evidence. It does not ask about tools, models, testing, or procedures. Those belong to the operators.

What you will have at the end

Forty-two answers, a one-page summary, a list of named owners, and a list of the places the organization answered "don't know." The last list is the most useful.

Govern the system, not the technology. The question is not whether any single AI model is safe. It is whether the organization has a program designed and operating well enough that its leaders can make an informed decision about the AI they are already using.

The oversight layer of AI governance

Where the duty sits

AI governance happens at three layers. At the technical layer, engineers and vendors test models, secure data, and build controls. At the operational layer, management writes the policies and runs the risk processes that direct that work. At the oversight layer, the Governing Body and executive management decide what must be governed, who is accountable for it, and what evidence they will require. Most frameworks, most consultancies, and most internal effort go to the bottom two layers. The duty sits at the top.

The Center's domain



OVERSIGHT GOVERNANCE

WHAT must be governed. WHO is accountable.

Governing Body and executive management. The duty of oversight sits here.
This review examines only this layer.

Left to the operators



OPERATIONAL GOVERNANCE

Policies, procedures, risk processes, monitoring.

Management and the second line. The program directs it; the Center does not prescribe it.

TECHNICAL GOVERNANCE

Models, data, testing, controls, tooling.

Engineering, the first line, and vendors. HOW it is done.

Under Delaware law the board owes a duty of oversight. *In re Caremark International Inc. Derivative Litigation* (Del. Ch. 1996) requires directors to put a reporting system in place for the corporation's mission critical risks. *Marchand v. Barnhill* (Del. 2019) held that the system must actually exist and operate for the risk in question, not for risk in general. As AI becomes material to an organization's business, it falls inside that duty. Regulators in financial services, healthcare, and public disclosure expect the same thing in their own terms: demonstrable oversight, not adoption alone.

The oversight layer is the Center's domain. The Center defines what boards and executives must govern and who is accountable. It does not prescribe how. This review is that definition applied at the domain level, in a form an organization can use on its own.

The program is the governance. A committee discussing AI is not a program. A set of model controls is not a program. A charter, decision rights, risk tolerances set in advance, escalation paths, a reporting cadence to the Governing Body, and a traceable evidentiary record are. That is what the fourteen domains that follow look for.

Five pillars, fourteen domains

The AI Oversight Program, at the level this review works

The AI Oversight Program organizes oversight governance into five pillars and fourteen domains. Beneath the domains sit seventy-six controls and the diagnostic statements the Center uses to assess maturity. This review stops at the domain. That is deliberate: the domain is the level at which a director or an executive can answer from knowledge rather than from a briefing.



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

Three domains matter more than the rest

Domain 01 creates the mandate for the program. Domain 02 gives it the authority and the resources to act. Domain 13 makes sure critical risks reach the people who can decide. A shortfall in any of the three is not one gap among many. Without them there is no program for the other eleven domains to belong to.

How to complete the review

Three questions, four answers, one rule

Every domain is put to the same three questions, worded for that domain. Together they are the test the Center applies to any oversight claim: what must be governed, who is accountable, and how it is evidenced.

Exists

Is the oversight element in place? Not planned, not discussed, not in draft. Adopted, approved, and current.

Accountable

Is one role, function, or body named as accountable for it? Write the role down. If two names come to mind, or none, the answer is not "present."

Evidenced

Could the organization produce the record if a regulator, an examiner, or a court asked for it this month? Minutes, approvals, reports, dated documents.

Four answers. **Present** means the element is in place and current. **Partial** means it exists in part, for some AI activities, or in draft. **Absent** means it does not exist. **Don't know** means the people completing the review cannot say. Treat "don't know" as a finding in its own right. An oversight element that the accountable executives cannot locate is, for oversight purposes, an element the organization does not have.

One rule. Answer at the oversight layer. If an answer begins with the name of a tool, a platform, a testing method, or a procedure, the answer has dropped a layer. Ask the question again: what did the Governing Body require, who did it hold accountable, and what record shows it.

Complete the review in one sitting with the people who would be accountable, then transfer the answers to the summary on page 12. The pattern across fourteen domains is the result; no single answer is.

Proportionality applies. A smaller organization with a modest AI footprint will meet several of these elements through existing risk committee structures and existing charters. That is acceptable. The test is whether the element exists and someone owns it, not whether it has its own document.

I Agile Governance *The Constitution*

The mandate, the structure, and the authority that give the program its standing. Without this pillar there is no program to review.

01 AI Governance Program and Policy Framework

The mandate for the program itself.

EXISTS	The Governing Body has adopted a written mandate for AI oversight, whether a charter, a resolution, or an amendment to an existing charter, that names AI as a matter of oversight.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	One executive or one constituted body is named as accountable for the AI oversight program as a whole.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	The mandate and the policy framework beneath it can be produced on request, with approval dates and the date of last review.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

02 AI Governance Structure, Oversight, and Resources

The authority and the resources to act.

EXISTS	The body that oversees AI has defined decision rights: what it approves, what it can stop, and what it must escalate.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	The accountable function has budget, staff, and a reporting line adequate to the mandate, and the Governing Body has confirmed that.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	Minutes or records show the body exercising those decision rights, not only receiving updates.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

03 Governance Program Assurance and Continuous Learning

The program reviews itself.

EXISTS	The oversight program is reviewed on a defined cycle, by internal audit or an independent party, as a subject of assurance in its own right.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	A named function owns the review cycle and the follow-up on its findings.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	The organization can show a change made to the program as a result of a review, a lesson learned, or a regulatory development.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

II Risk-Informed System *The Guardrails*

The appetite, the scope, and the intelligence that let the organization tell an acceptable AI risk from an unacceptable one before the decision is made.

04 AI Risk Methodology, Scope, and Tolerance

What is in scope, and how much risk is acceptable.

EXISTS	The Governing Body has approved a statement of AI risk appetite specific enough that management can tell when an activity exceeds it. A general commitment to responsible AI does not meet this test.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	A named function owns the risk methodology and the scope decision: which AI activities the program reaches, including third-party AI, AI embedded in enterprise products, employee use of generative AI, and agentic systems.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	Tolerances are recorded, dated, and have been revisited as the organization's use of AI has changed.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

05 Risk Intelligence and Threat Landscape

Knowing what is changing outside the organization.

EXISTS	The program has a defined means of learning about emerging AI risks, regulatory developments, and incidents elsewhere, and of deciding which of them matter here.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	A named function is accountable for bringing that intelligence to the oversight body on a cadence.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	The oversight body can point to a decision or a change it made in response to external risk intelligence.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

III AI Trust and Assurance *The Evidence*

The standards AI systems must demonstrably meet, and the assurance that reaches the oversight body to prove they meet them.

06 AI Model Risk and Agentic Lifecycle Oversight

Standards for AI systems across their life, including systems that act.

EXISTS	The Governing Body has set the standards that AI systems, including agentic systems, must demonstrably meet before deployment and while in use, and has required a current inventory of them.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	Ownership is assigned for every AI system in the inventory, and a function is named to confirm that the inventory is complete.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	The oversight body receives attested reporting, on a defined cadence, on whether systems meet the standards, including material degradation and incidents.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

07 AI Data Governance Oversight

What the organization requires of the data AI learns from and acts on.

EXISTS	The program states what is required of the data behind AI systems: its provenance, its quality, its lawful use, and its protection.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	A named function is accountable for data governance as it applies to AI, distinct from the teams that build the systems.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	The oversight body can see where data requirements are not met and what has been decided about it.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

III AI Trust and Assurance *The Evidence (continued)*

08 AI Transparency, Explainability, and Human Oversight

When a decision must be explained, and when a person must be able to intervene.

EXISTS	The program states when an AI-influenced decision must be explainable, to whom, and when a person must be able to review, intervene, or reverse it, in proportion to the risk.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	For each higher-risk use, a named role is accountable for human oversight of the system's outputs and actions.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	For a higher-risk system, the organization can show the record of human review and any intervention that occurred.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

09 AI Security and Resilience Assurance

Assurance that AI systems are secured and that the organization can carry on if one fails.

EXISTS	The program requires assurance that AI systems are protected against misuse, manipulation, and failure, and that the organization can continue to operate if a material AI system fails or must be withdrawn.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	Accountability for AI security and resilience is assigned and reconciled with the existing security and business continuity functions, so that nothing falls between them.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	Results of independent assurance reach the oversight body, with the material findings and their disposition.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

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

The link between oversight and the decisions that commit capital, accept risk, and bring AI in from outside.

10 Risk-Informed Strategy, Resources, and Organizational Readiness

Oversight before the commitment, not after it.

EXISTS	AI oversight has a place in the organization's strategy and investment decisions before commitments are made, not only at deployment.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	A named executive is accountable for confirming that the organization has the capacity, skills, and resources to govern the AI it plans to adopt.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	Records of major AI investment decisions show that oversight considerations were presented and weighed.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE ROLE		EVIDENCE NOTED

11 AI Value Realization and Operational Resilience Oversight

Whether AI is delivering what was approved, within the risk that was accepted.

EXISTS	The program requires that the expected value and the accepted risk of material AI initiatives be stated at approval and reported against afterward.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	A named function owns reporting on whether AI initiatives are delivering what was approved and operating within the risk that was accepted.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	The oversight body has received at least one such report and can show what it decided as a result.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE ROLE		EVIDENCE NOTED

12 Third-Party AI and Supply Chain Governance

AI the organization consumes rather than builds.

EXISTS	The program treats AI supplied by vendors, embedded in enterprise products, and consumed from foundation model providers as within scope, with requirements set for each.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	Accountability for AI risk that originates with a supplier is assigned inside the organization, not left with the supplier.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	The inventory identifies third-party AI dependencies, and the oversight body has visibility into the material ones.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE ROLE		EVIDENCE NOTED

V Risk Escalation and Disclosure *The Voice*

How a critical risk reaches the people who can decide, how the organization discloses it, and how it knows the path works.

13 AI Risk Escalation and Disclosure Protocols

Critical risks bypass silos and reach decision makers.

EXISTS	Escalation thresholds for AI risks and incidents are defined, from the first line to the second and from management to the Governing Body, with timing, and the organization's AI disclosure obligations to regulators, customers, and the market are identified.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	A named role decides, in a specific event, whether a threshold has been met and makes the escalation, and a named role owns each disclosure obligation.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	The escalation path has been used or tested, and the organization can show a case that reached the level the protocol required.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

14 Validation of Escalation and Governance Effectiveness

Testing the emergency brakes.

EXISTS	The program tests whether escalation and oversight actually work: that risk information reaches decision makers intact, in time, and in a form that allows a decision.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
ACCOUNTABLE	A function independent of the first line is accountable for that testing and for reporting the results to the Governing Body.	☐ Present ☐ Partial ☐ Absent ☐ Don't know
EVIDENCED	Within the last reporting cycle, the Governing Body has received a report on the effectiveness of escalation and oversight, not only a summary of incidents.	☐ Present ☐ Partial ☐ Absent ☐ Don't know

ACCOUNTABLE ROLE

EVIDENCE NOTED

SUMMARY

Your baseline at a glance

Write P, Pa, A, or DK in each cell. Read the pattern, not the count.

Domain	Exists	Accountable	Evidenced
I Agile Governance · <i>The Constitution</i>			
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 · <i>The Guardrails</i>			
04 AI Risk Methodology, Scope, and Tolerance			
05 Risk Intelligence and Threat Landscape			
III AI Trust and Assurance · <i>The Evidence</i>			
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 · <i>The Strategic Alignment</i>			
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 · <i>The Voice</i>			
13 AI Risk Escalation and Disclosure Protocols			
14 Validation of Escalation and Governance Effectiveness			

P = Present · Pa = Partial · A = Absent · DK = Don't know. There is no total. Fourteen "present" answers with no named owners is a weaker result than nine "present" answers with an owner in every domain.

Reading your baseline

Five patterns the summary will show

Owners named in every domain, evidence in most

A program exists. The next question is how deeply it is institutionalized: whether it operates the same way when the people change, whether it is measured, and whether it improves. That is a maturity question, and it takes evidence to answer.

"Present" under Exists, "Absent" or "Don't know" under Accountable

Repeated across domains, this is the Committee Fallacy. AI is being discussed, perhaps by a committee, and no one owns the outcome. A committee is an oversight body; it needs a program to oversee.

Strong in Pillar III, weak in Pillars I, IV, and V

The Assurance Trap. The organization is testing its models while the executive decisions that set the risk, and the escalation paths that carry it upward, go ungoverned. Model controls are necessary. They are not oversight.

"Don't know" in Domains 01, 02, or 13

Treat these before anything else. Without a mandate there is no program. Without authority and resources nothing can be acted on. Without escalation, every other failure has nowhere to go.

Answers that describe tools, testing, or procedures

The review has been completed at the operational layer. The elements may well exist below the oversight layer, but the Governing Body cannot demonstrate that it required them, owns them, or sees their results. Ask the questions again with the board's evidence in mind.

FROM BASELINE TO MATURITY

The baseline tells the organization whether the oversight layer exists. The Center's AI Governance Maturity Model tells it how far along that layer is. The Model assesses seventy-six controls across the same fourteen domains at five maturity levels, requires traceable evidence for every score, and checks for the patterns that produce the appearance of oversight without its substance. Assessments are conducted through the Center's AI Oversight Governance Portal, in facilitated engagements, and through the Center's channel partners in the accounting and legal professions. The Center does not sell technology, recommend vendors, or perform implementation work.

Center for AI Oversight

Defining oversight governance for AI in regulated industries

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. It leaves operational design to the operators. The Center is guided by former senior leaders from the Federal Reserve, the Securities and Exchange Commission, the Federal Communications Commission, the U.S. Department of the Treasury, and the Department of Health and Human Services, and is a contributing member of the NIST AI Safety Institute Consortium.

The Center's work takes four forms: the AI 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 with institutional partners; education for directors, audit executives, and risk leaders; and convening and speaking in boardrooms, at conferences, and in regulator forums.

Terms of use

This review is published as an educational resource. An organization may reproduce and use it for its own governance purposes, and advisers may use it with their clients, with attribution to the Center for AI Oversight. It may not be sold, incorporated into a commercial product or service, or presented as a maturity assessment or as an assessment conducted by the Center. It does not constitute legal advice.

Discuss your baseline

The Center offers a short conversation on the results of a completed review, without charge and without obligation. Write to info@cfaio.org with the summary page. Organizations that want a full maturity assessment can begin there.



Center for
AI Oversight

Build Trust. Enable Speed. Lead Responsibly.

cfaio.org · info@cfaio.org

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

© 2026 Center for AI Oversight. All rights reserved. First edition · September 2026.