CENTER FOR AI OVERSIGHT

The AI Oversight Program Program Guide

Oversight governance that enables Decision Velocity

What boards and executives must govern in AI, and who is accountable.
The five pillars and the fourteen domains, explained.

Center for AI Oversight · A 501(c)(3) nonprofit · cfaio.org · Version 1.0, November 2026

ABOUT THIS GUIDE

Terms of use

Read it, share it, cite it. Do not sell it, alter it, or present it as an assessment.

The AI Oversight Program: Program Guide, version 1.0, November 2026. Published by the Center for AI Oversight, a 501(c)(3) nonprofit educational institution. cfaio.org.

This Guide is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International license (CC BY-NC-ND 4.0). You may read, share, and cite it, with attribution to the Center for AI Oversight, for internal governance work. You may not sell it, incorporate it into a commercial product or service, alter it, or present it as a maturity assessment.

Advisers may draw on this Guide in their own work with clients, with attribution, provided the Guide itself is not sold, rebranded, or altered.

The controls, diagnostic statements, maturity levels, and scoring logic of the AI Oversight Program are proprietary to the Center. They are not reproduced here and are not licensed by this document.

The Guide describes the oversight layer of AI governance. It defines what must be governed and who is accountable. It does not prescribe how, and nothing in it is legal advice.

Requests for other uses, and questions about the Baseline Review or a maturity assessment: info@cfaio.org.

© 2026 Center for AI Oversight. All rights reserved except as licensed above.

FOREWORD

Why this Guide exists

The deck makes the case. This Guide is the reference.

The Center for AI Oversight publishes a short overview of the AI Oversight Program for directors and executives who need to understand it in a single sitting. This Guide is for the reader who has to go further: the audit executive deciding whether the program belongs in the audit plan, the second-line risk leader deciding what to ask of management, the general counsel deciding whether the structure would hold up when the duty of oversight is tested, and the examiner deciding whether it is credible.

Those readers do not need more persuasion. They need to know what each pillar and each domain requires, who is accountable for it, and where one domain ends and the next begins. That is what this Guide gives them.

One distinction runs through every page. AI governance has three layers: oversight sets the program, management runs it, and technical controls enforce it. The Program is the oversight layer. It is not a management framework and it is not a control catalogue, and a committee that reviews technical controls is not doing oversight. Most of what is called AI governance today flattens those three layers into one. This Guide does not.

Two things it does not do. It does not prescribe how any of this is done. The Program defines what must be governed and who is accountable, and leaves operational design to the people who run the organization. And it does not reproduce the controls, the diagnostic statements, or the maturity scoring. Those are the instrument behind the Program, and they are used in an assessment, not read in a guide.

The Program exists because the duty of oversight already reaches AI and most organizations cannot yet show they have discharged it. Closing that gap is not a matter of slowing down. It is a matter of knowing what the duty requires, assigning it to someone, and being able to prove it was met. The organizations that do this will make decisions about AI faster than those that do not, and with evidence behind them.

Brian J. Allen, Esq.

Executive Director, Center for AI Oversight

CONTENTS

In this Guide

Three parts: the case, the architecture, the Center.

Terms	Terms of use	2
Foreword	Why this Guide exists	3
Part One		
1	Why the Program exists	5
2	What the Program defines	7
Part Two		
3	The architecture	10
Pillar I	Agile Governance	12
Pillar II	Risk-Informed System	14
Pillar III	AI Trust and Assurance	15
Pillar IV	Risk-Based Strategy and Execution	17
Pillar V	Escalation and Disclosure	19
9	Builders and Buyers	20
10	The maturity model	21
Part Three		
11	Built on converging authority	23
12	The Center for AI Oversight	24

PART ONE · CHAPTER 1

Why the Program exists

AI is being adopted faster than most organizations can govern it. The Program is the institutional capability that closes that gap.

The Velocity Gap

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. It widens over time unless something closes it. The Program is that something. The organizations that govern AI fastest will outperform those that govern it best on paper.

Governance is not a brake on AI adoption. It is what makes confident adoption possible. Risk tolerances set in advance become decision points, not refusals. A risk identified is information that enables a decision, not an obstacle that stops one.

The duty of oversight is already settled

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. Three things follow.

- **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.

What organizations need from a Program

Informed decisions

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

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.

A different speed

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

Defensible governance

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

What looks like oversight, but is not

Three patterns reproduce the appearance of oversight without its substance. The Program is built to prevent all three.

- **The Committee Fallacy.** A committee discussing AI is not a program. A committee is an oversight body, and it needs a functioning program to oversee.
- **The Assurance Trap.** Auditing model controls while ignoring the executive decisions that set risk. Technical assurance over a hollow decision structure is compliance theatre.
- **Framework Fatigue.** A patchwork of standards with nothing internal to anchor them. Without an internal program, external standards become checklists.

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; and a traceable evidentiary record. The program is the governance. It is what makes oversight defensible when the duty is tested.

PART ONE · CHAPTER 2

What the Program defines

Govern the system, not the technology. The Program defines WHAT must be governed and WHO is accountable. It never defines HOW.

The oversight layer

AI governance has three layers. **Oversight** sets the program: the mandate, the appetite, the standards, the escalation paths, and the evidence the Governing Body requires. **Management** runs it: the decisions, the processes, and the people inside the tolerances oversight set. **Technical controls** enforce it: the model, data, and security measures that make the standards real in the systems. Defined exchanges flow down and up between the layers.

The Program is the oversight layer, and only that layer. It does not replace a management framework or a control catalogue, and it is not satisfied by either. Flattening the three into one is how a committee comes to believe that reviewing technical controls is oversight. Everything in this Guide sits in the first layer.

WHAT we define, and 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. Throughout this Guide, the Governing Body requires, approves, confirms, sees, and receives. Management defines, sets, runs, and decides inside tolerance. That line is deliberate. Heads in, hands off.

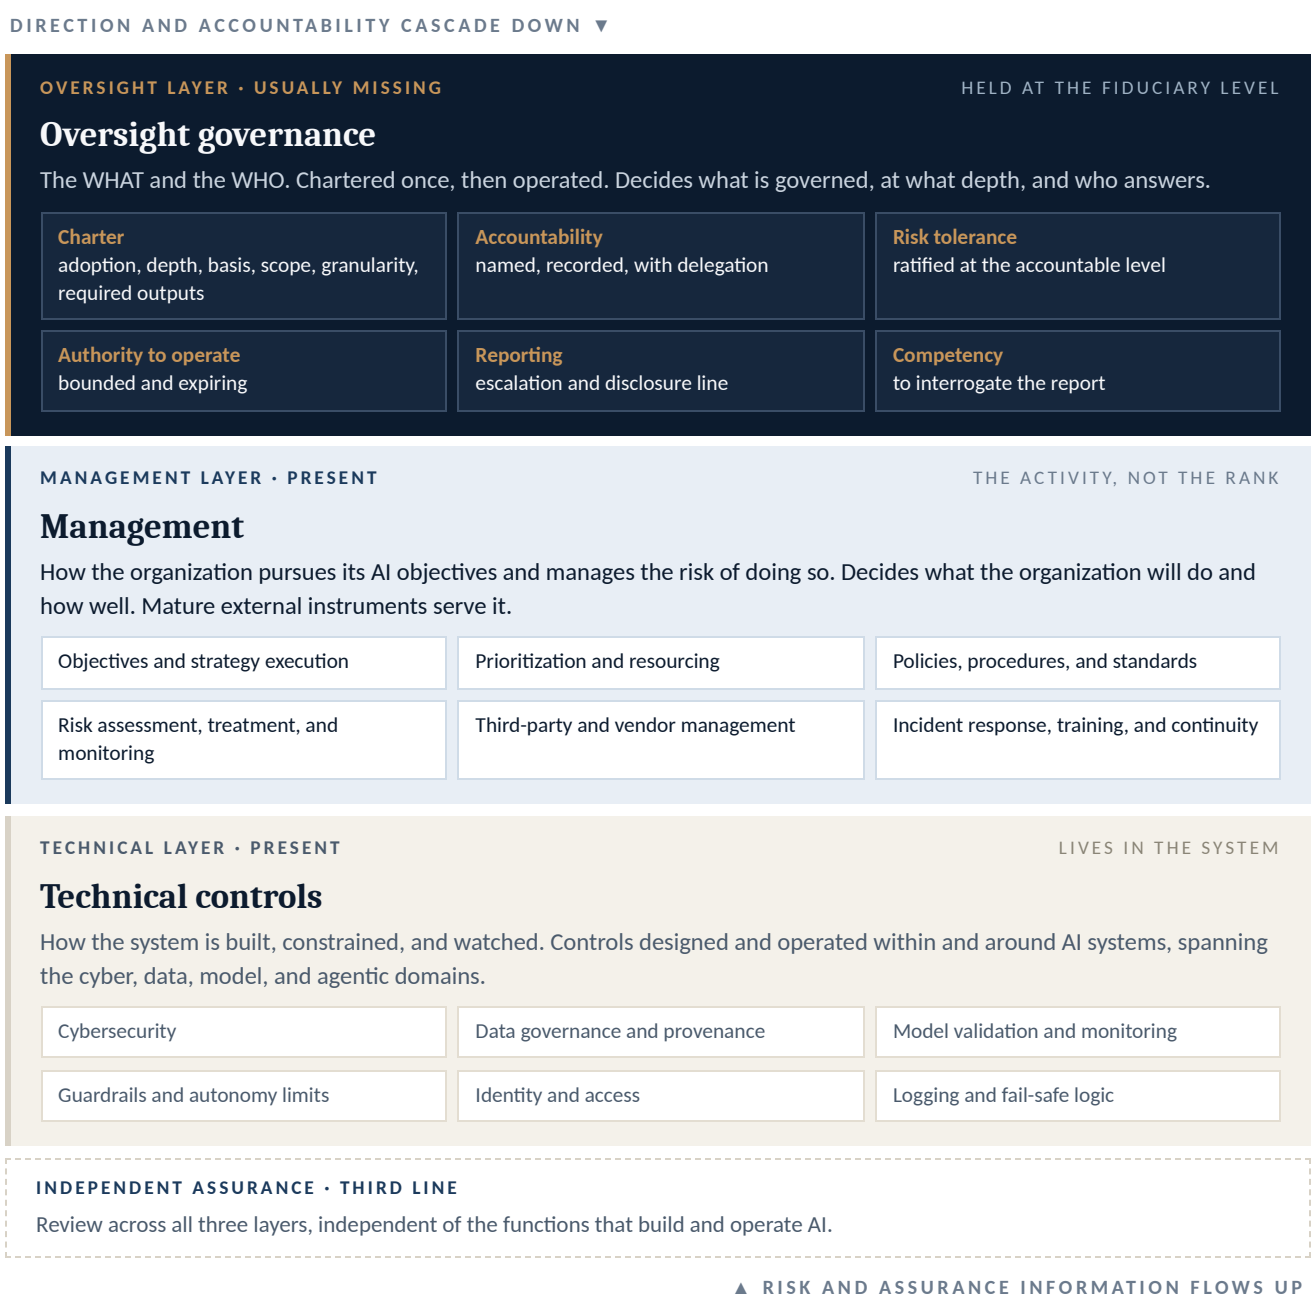


FIGURE 1 · THE LAYERS OF AI GOVERNANCE. Layer assignment follows subject matter, not audience altitude. Management names the activity, not the rank: when the chief executive sets an AI objective, that is management; when the same executive names who answers for a deployment, that is oversight. The Program is the top layer. From Center Paper No. 7, *The Oversight Layer of AI Governance* (September 2026).

The convictions the Center is built on

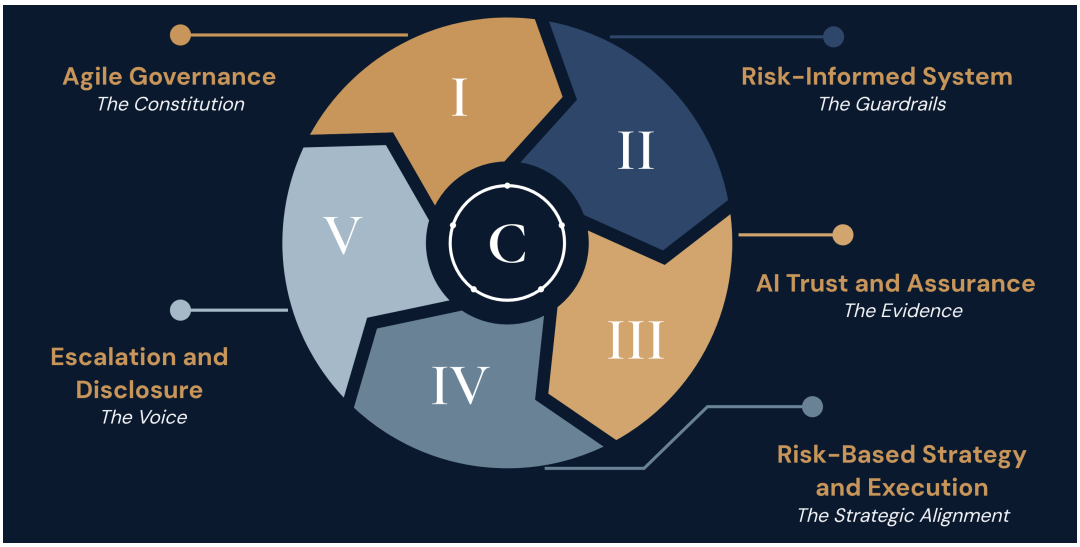
The Bedrock Principles, in brief.

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.

PART TWO · CHAPTER 3

The architecture

Five pillars, fourteen domains, seventy-six controls, three hundred seven diagnostic statements. One defensible standard of care.



Each pillar carries a constitutional role: what it is to the program, in one phrase. Each pillar is made up of its domains, and each domain states what must be governed and who is accountable. Beneath the domains sit the controls and the diagnostic statements that make up the maturity assessment. This Guide stops at the domain layer. The pillars are locked across every version of the Program; the domains and controls evolve beneath them.

	Pillar	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	**Escalation and Disclosure** The Voice	13 AI Risk Escalation and Disclosure Protocols 14 Validation of Escalation and Governance Effectiveness

How to read the domain chapters

Each domain in the chapters that follow is set out the same way. The opening line is the domain's definition. The paragraph states what must be governed and who is accountable. The board question is the question a director can put to management that tells, in one answer, whether the domain is in place. The closing lines mark where the domain ends and its nearest neighbor begins, so that the fourteen can be read as complete and non-overlapping. Where the obligation differs for an organization that deploys AI built by others, a Buyer note says how.

PART TWO · PILLAR 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.

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

01 AI Governance Program and Policy Framework

The mandate for the program itself. This domain governs whether AI oversight exists as a matter of record: a charter, a resolution, or an amendment to an existing charter that names AI as something the board oversees, and the policy set that sits beneath it. The mandate also fixes the program's scope and grain: which parts of the organization and which kinds of AI use it reaches, and at what level of detail, sized to this organization rather than copied from a template. That scope is defined, documented, and visible to the board. The Governing Body owns the mandate. One executive or one constituted body is named as accountable for the program as a whole.

> **THE BOARD QUESTION**
> Can we produce the document that says we govern AI, what it covers and at what grain, and the date we last reviewed it?

Domain 01 establishes that the program exists, what it reaches, and who owns it. Domain 02 establishes what that owner is allowed to decide and what they have to decide with.

02 AI Governance Structure, Oversight, and Resources

The authority and the resources to act. This domain governs the decision rights of the body that oversees AI: what it approves, what it can stop, and what it must escalate. It also governs whether that body is funded and staffed to the size of its mandate, and whether anyone has assessed that fit rather than assumed it. The accountable function owns the structure and the reporting line. The Governing Body confirms that the budget and staff match the mandate, because an oversight body that can only receive updates is not an oversight body.

> **THE BOARD QUESTION**
> How does management assess whether the oversight function has the resources to develop and adopt what the mandate requires, and has the board seen that assessment?

Domain 02 is about the standing structure. Domain 13 is about what happens inside that structure when a specific risk crosses a threshold.

03 Governance Program Assurance and Continuous Learning

The program reviews itself, as designed. This domain governs whether the oversight program is a subject of assurance in its own right: reviewed on a defined cycle by internal audit or an independent party, and changed as a result of what the review finds, what the organization learns from incidents, and what regulators publish. A named function owns the review cycle and the follow-up.

> **THE BOARD QUESTION**
>
> Has the AI oversight program itself been audited, or is it on the audit plan, and is the cadence right for how fast our use of AI is changing?

Domain 03 tests the program as it was designed. Domain 14 tests whether the escalation path works in use. One asks whether the structure is sound; the other asks whether the brakes engage.

PART TWO · PILLAR 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. The pillar turns on whether a defined methodology exists for setting tolerance and identifying risk. The Governing Body requires that methodology and approves the tolerance. The decisions made inside it belong to management.

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

04 AI Risk Methodology, Scope, and Tolerance

What is in scope, and how much risk is acceptable. This domain governs the methodology every other domain depends on: how AI risk is identified, how tolerance is set, which AI activities the program reaches, and whether that methodology has been approved as appropriate for this organization. Scope has to name the hard cases: third-party AI, AI embedded in enterprise products, employee use of generative AI, and systems that act on their own. Tolerance has to be specific enough that management can tell when an activity exceeds it. A general commitment to responsible AI does not meet that test. The Governing Body approves the appetite and confirms the methodology fits the organization. A named function owns the methodology, the scope decision, and the running of use cases through it.

> **THE BOARD QUESTION**
> Are the AI use cases that carry risk being worked through the methodology, and which ones are not?

Domain 04 sets the boundary. Domain 06 holds individual systems to the standard inside it.

05 Risk Intelligence and Threat Landscape

Knowing what is changing outside the organization. This domain governs whether a standing capability exists to learn about emerging AI risks, regulatory developments, and incidents elsewhere, and to decide which of them matter here. The point is not a news feed. The point is a defined intelligence function with an owner, a method for deciding relevance, and a cadence for reaching the oversight body. A named function is accountable for it.

> **THE BOARD QUESTION**
> Is there a standing AI threat and regulatory intelligence capability, with an owner and a cadence to the oversight body, or do we learn about outside events by accident?

Domain 05 looks outward at what is changing. Domain 09 looks inward at whether our own systems are protected against it.

PART TWO · PILLAR 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.

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

06 AI Model Risk and Agentic Lifecycle Oversight

Standards for AI systems across their life, including systems that act. This domain governs what every AI system in scope must demonstrably meet before it is deployed and while it is in use, and it requires a current inventory of those systems so that the standard has something to apply to. Agentic systems, which take actions rather than produce outputs, are inside this domain, not a special case outside it. The Governing Body requires that the standards exist and that they are approved. Management sets them and holds systems to them. Ownership is assigned for every system in the inventory, and a function is named to confirm the inventory is complete. The oversight body receives attested reporting on whether systems meet the standard, including degradation and incidents.

> **THE BOARD QUESTION**
>
> How many AI systems do we run, who owns each one, and which ones are not meeting the standard today?

Buyers: the standard applies to the systems you deploy, not only the ones you build, and the inventory of them is yours to keep.

Domain 06 governs every system in scope across its life. Domain 12 governs the terms on which outside systems enter that scope.

07 AI Data Governance Oversight

What the organization requires of the data AI learns from and acts on. This domain governs the requirements placed on the data behind AI systems: where it came from, whether it is fit for the purpose, whether its use is lawful, and how it is protected. It does not duplicate the organization's existing data governance. It states what that governance has to cover once AI is learning from the data and acting on it. A named function is accountable for data governance as it applies to AI, separate from the teams that build the systems.

> **THE BOARD QUESTION**
> Do we know what data our AI systems were built on, and whether we had the right to use it?

Buyers: the requirement reaches the data you feed a vendor system and the data its outputs act on.

Domain 07 governs the data. Domain 06 governs the system that uses it.

08 AI Transparency, Explainability, and Human Oversight

When a decision must be explained, and when a person must be able to intervene. This domain governs the organization's own rules for AI-influenced decisions: which decisions must be explainable, to whom, and when a person must be able to review, intervene, or reverse the outcome. The rule is proportionate to the risk; not every decision needs a human in the loop, and the program says which do. Where it does, the human role is a documented procedure, not a habit. Explainable means the decision process and the information behind it can be set out; it does not require that the model itself be interpretable. For each higher-risk use, a named role is accountable for human oversight of the system's outputs and actions.

> **THE BOARD QUESTION**
> Are our high-risk AI decision processes documented, with the human role defined as a procedure, and can each decision be explained even where the model cannot be interpreted?

Domain 08 governs the human role in decisions the system makes. Domain 13 governs what happens when the system, or the human, finds a risk that has to go up.

09 AI Security and Resilience Assurance

Assurance that AI systems are secured and that the organization can carry on if one fails. This domain governs two requirements: that AI systems are protected against misuse, manipulation, and failure, and that the organization can keep operating if a material AI system fails or must be withdrawn. The oversight concern is the seam. AI security and AI resilience sit between the existing security function and the existing business continuity function, and this domain requires that accountability be assigned and reconciled so that nothing falls between them. Resilience is measured, not assumed: indicators exist for the material systems, with thresholds, and they move as the AI environment moves. Independent assurance results reach the oversight body with the material findings and their disposition.

> **THE BOARD QUESTION**
> Are resilience indicators for our material AI systems established, with thresholds, and adjusted at the pace the AI environment changes?

Domain 09 is about whether the system survives attack or failure. Domain 11 is about whether the business gets what it was promised from it.

PART TWO · PILLAR 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. This is where the balance of risk and reward is struck, and where it is revisited at the pace this environment moves. The board sees the balance. Management strikes it.

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

10 Risk-Informed Strategy, Resources, and Organizational Readiness

Oversight before the commitment, not after it. This domain governs whether AI oversight has a seat at the point where strategy and investment decisions are made, rather than arriving at deployment when the money is spent and the direction is set. It also governs readiness: whether the organization has the capacity, skills, and resources to govern the AI it plans to adopt, before it adopts it. A named executive is accountable for confirming that readiness. Records of major AI investment decisions show that oversight considerations were presented and weighed.

> **THE BOARD QUESTION**
>
> Taken together, are our AI investments inside the appetite and tolerance we set for risk and reward, and can the board see that fit without stepping into the investment decisions themselves?

Domain 10 governs the decision to commit. Domain 11 governs what happened after the commitment was made.

11 AI Value Realization and Operational Resilience Oversight

Whether AI is delivering what was approved, within the risk that was accepted. This domain governs the closing of the loop. At approval, the expected value and the accepted risk of a material AI initiative are stated. Afterward, the program requires reporting against both, so the oversight body can see whether the initiative delivered and whether it stayed inside the risk that was agreed. A named function owns that reporting. The oversight body has received at least one such report and can show what it decided as a result.

> **THE BOARD QUESTION**
>
> Of the AI initiatives we approved, which delivered, which did not, and what did we do about the ones that did not?

Domain 11 governs performance against the approval. Domain 09 governs whether the system survives attack or failure.

12 Third-Party AI and Supply Chain Governance

AI the organization consumes rather than builds. Most organizations do not build the AI they rely on. They license it, inherit it inside vendor platforms, or consume it from foundation model providers. This domain governs all of it as in scope, with requirements set for each kind, and it requires that the organization understand what a supplier failure or withdrawal would mean for its own operations and whether that exposure fits the risk and reward it accepted. Accountability for AI risk that originates with a supplier is assigned inside the organization. A vendor can be the source of the risk. It cannot be the owner of it. The inventory identifies third-party AI dependencies, and the oversight body has visibility into the material ones.

> **THE BOARD QUESTION**
> Do we understand the resilience implications of the AI we depend on from outside, does that exposure fit the risk and reward profile we accepted, and who inside is accountable for knowing?

Buyers: this is where most of your AI risk originates. Accountability for it stays inside the organization, not with the supplier.

Domain 12 governs the terms on which outside AI enters scope. Domain 06 governs every system, outside or inside, once it is in.

PART TWO · PILLAR V

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. A tolerance is not a no. It is a decision point. When a risk falls outside tolerance, the question is whether the decision reaches the people with the authority to make it, in time, and whether the organization meets its disclosure responsibilities once it does.

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

13 AI Risk Escalation and Disclosure Protocols

Critical risks bypass silos and reach decision makers. This domain governs the path a serious AI risk or incident takes upward: the thresholds that trigger escalation from the first line to the second and from management to the Governing Body, the timing attached to each, and the speed at which a material risk is recognized as material in the first place. It also governs disclosure, which arises for three reasons: because a risk is material and disclosure is an obligation, because it meets a threshold a regulator demands, or because the public's trust depends on hearing it from the organization first. A named role decides, in a specific event, whether a threshold has been met and makes the escalation. A named role owns each disclosure obligation.

> **THE BOARD QUESTION**
> When a material AI risk emerges, how fast do we know, and do we know which of our disclosure obligations it triggers?

Domain 13 governs the path up. Domain 14 governs whether the path works.

14 Validation of Escalation and Governance Effectiveness

Testing the emergency brakes, as used. This domain governs whether escalation and oversight actually function: whether risk information reaches decision makers intact, in time, and in a form that allows a decision. The test is independent of the first line, and its results go to the Governing Body as a report on effectiveness, not a summary of incidents. A function independent of the first line is accountable for the testing and the reporting.

> **THE BOARD QUESTION**
> Have the AI risk escalation protocols been tested, and did they work?

Domain 14 tests the escalation path in use. Domain 03 reviews the program as designed. Both are assurance; they look at different things.

PART TWO · CHAPTER 9

Builders and Buyers

Most organizations do not build the AI they rely on. The Program governs them anyway, and it says how the obligation differs.

The Program distinguishes two kinds of organization. A **Builder** develops AI systems: it trains models, assembles agents, or engineers products on top of foundation models, and it owns the design decisions that shape how the system behaves. A **Buyer** deploys AI systems built by others: licensed platforms, AI embedded in enterprise software, and capabilities consumed from foundation model providers. Most organizations in regulated industries are Buyers, and many are both at once, building in one part of the business while buying in the rest.

The duty is the same. The evidence differs.

Being a Buyer does not shrink the duty of oversight. The board still owes it, management still discharges it, and all fourteen domains still apply. What changes is where the evidence comes from and where the risk originates. A Builder's risk is concentrated in what it makes. A Buyer's risk is concentrated in what it brings in, in the data it feeds a vendor's system, and in what the outputs of that system are allowed to act on.

Three domains carry the largest difference, and each is marked with a Buyer note in the chapters above.

- **Domain 06, AI Model Risk and Agentic Lifecycle Oversight.** The standard applies to the systems you deploy, not only the ones you build, and the inventory of them is yours to keep. A Buyer that cannot list the AI systems it runs has not met this domain, whoever built them.
- **Domain 07, AI Data Governance Oversight.** The requirement reaches the data you feed a vendor system and the data its outputs act on. The vendor's own data practices are the vendor's; what your organization puts in and what it does with what comes out are yours.
- **Domain 12, Third-Party AI and Supply Chain Governance.** This is where most of a Buyer's AI risk originates. Accountability for it stays inside the organization, not with the supplier. A vendor can be the source of the risk. It cannot be the owner of it.

What this means for an assessment

The maturity assessment reads the same fourteen domains for a Builder and a Buyer. It differs in which diagnostic statements apply and in what counts as evidence. A Buyer is not asked to produce model validation records for a system it did not build. It is asked to show that it set a standard for that system, knows who owns it, holds the supplier to it, and can carry on if the supplier fails. That is a different body of evidence for the same obligation, not a lower bar.

PART TWO · CHAPTER 10

The maturity model

The domains say what must be governed. The maturity model measures how far it has been institutionalized.

Beneath the fourteen domains sit seventy-six controls, and beneath the controls, three hundred seven diagnostic statements. Together they form the instrument the Center uses to assess an organization's AI oversight maturity. This Guide describes the instrument. It does not reproduce it. The control text, the diagnostic statements, and the scoring logic are proprietary to the Center and are used in an assessment, not read in a document.

What the model measures

Maturity in this model means institutionalization, not automation. A high score does not mean the organization has the most tooling. It means the obligations of each domain are assigned, documented, exercised, evidenced, and reviewed, so that the oversight would survive the departure of the people who built it. The model assesses each control against five maturity levels, and every score traces to evidence the organization produced. Assertion does not score. A record does.

How a score is built

- **Controls.** Each domain resolves into controls: the specific obligations that, together, satisfy the domain. Every control cites the standards, regulations, and precedent that define the standard of care it reflects.
- **Diagnostic statements.** Each control resolves into diagnostic statements: the observable conditions that distinguish one maturity level from the next. These are what an assessment actually tests.
- **Evidence.** A diagnostic statement is met only when the organization can produce evidence that it is. The assessment records the evidence beside the score, so that the board can see not only where it stands but why.
- **Anti-pattern checks.** At the higher maturity levels, the model also checks for the patterns that reproduce the appearance of oversight without its substance: the Committee Fallacy, the Assurance Trap, and Framework Fatigue. An organization that exhibits one does not score as mature in the domain it affects, whatever else is in place.

What the result gives the board

A maturity assessment produces a board-grade report: where the organization stands on each pillar and domain, what evidence that rests on, which gaps matter most, and what an informed decision about the next step would weigh. It is written for the Governing Body and the executives accountable to it, in the language of

oversight rather than engineering. It answers one question above all: is the program designed and operating effectively, so that the organization can make informed decisions on the balance of risk and reward?

The free first step

Before an assessment, the Center offers the AI Oversight Baseline Review: a free, unscored self-review that asks three questions of each domain, whether it exists, whether it is owned, and whether it is evidenced. It stays at the domain level, gives away no control text, and tells an organization where to look first. The Baseline Review is available at cfaio.org.

PART THREE · CHAPTER 11

Built on converging authority

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

Technical standards

NIST AI RMF · ISO/IEC 42001 · ISO/IEC 38507

Repeatable risk processes and a voluntary governance structure.

Judicial precedent

Caremark · *Marchand* · Boeing

Boards must maintain functioning oversight systems for material risks.

Legislation and regulation

EU AI Act · SEC · FTC · state laws

Legally binding obligations and risk-tiered classification.

The AI Oversight Program

The interoperability layer.

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

None of these sources, on its own, tells a board what to govern. The standards describe processes. The regulations impose obligations. The precedent sets the duty. The Program is the layer that turns the three into one set of obligations a board can assign, evidence, and defend. Each control in the instrument cites the sources it rests on, so that an auditor or examiner can trace any score back to the authority behind it.

The outcome: Decision Velocity

Decision Velocity is 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.

PART THREE · CHAPTER 12

The Center for AI Oversight

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

The Center for AI Oversight defines oversight governance for AI: what boards and executives must govern, and who is accountable. It is an educational institution in the tradition of independent research and policy bodies. It defines what must be governed and who is accountable. It does not prescribe how, leaving operational design to operators.

How the Center engages

1. **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.
2. **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.
3. **Education and training.** NASBA-compliant education for directors, audit executives, and risk leaders. Governance competency for senior professionals, not technology tutorials.
4. **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. The Center is 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 and 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

Using this Guide

This Guide may be used freely for internal governance work and by advisers with attribution to the Center for AI Oversight. It may not be resold, incorporated into commercial products, or presented as a maturity assessment. The controls, diagnostic statements, and scoring logic of the AI Oversight Program are proprietary to the Center and are not reproduced here. To discuss an assessment, the Baseline Review, or the Center's work, write to info@cfaio.org.

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

© 2026 Center for AI Oversight. All rights reserved. The AI Oversight Program: Program Guide, version 1.0.