The Center for
AI Oversight

CENTER PAPER NO. 4 · JUNE 2026

Builders and Buyers

Oversight obligations that do not transfer. Board oversight of AI the institution bought rather than built, and why the contract does not move the accountability.

By **Brian J. Allen**, Executive Director

> *Each section moves in three tiers: a **board line**, a **why it matters** note, and a shaded **detail** block carrying the authority.*

EXECUTIVE SUMMARY

Accountability stays home

KEY TAKEAWAYS

1. Builders develop AI; Buyers deploy AI built by others. Most regulated institutions are Buyers, and the dominant governance failure of 2026 is the assumption that vendor diligence discharges the oversight duty. It does not.
2. Four bodies of law, built by different regulators for different industries, have arrived independently at one principle: the institution deploying AI owns the risk, whoever built it.
3. The Buyer governs across an information boundary it does not control. Buyer-side oversight is its own discipline, not Builder-side oversight applied with less information.
4. The Buyer's program needs four elements: a vendor-embedded inventory, named ownership of bought systems, escalation that reaches through the vendor relationship, and a record kept with built-AI discipline.
5. The Buyer's obligations are no lighter than the Builder's, and heavier in one respect: the Buyer answers for a system it did not design and cannot fully inspect. The program is the governance.

This paper makes its argument by demonstration rather than assertion. When four unrelated regimes reach the same conclusion without coordinating, the conclusion belongs to the oversight layer, not to any sector. The Builder statutes themselves now subsidize the Buyer's work, because the safety frameworks and incident reports they mandate are diligence material the law produced for exactly this purpose.

PART ONE

The segmentation, and why it changes the question

Builders and Buyers is a segmentation of oversight obligation, not a vendor-versus-customer relabeling. The Builder's duties attach to what it creates; the Buyer's attach to what it does with what it acquired.

Why it matters. The regimes now write the split explicitly: the EU AI Act allocates duties between providers and deployers, Colorado's replacement statute between developers and deployers, and the employment statutes bind employers regardless of who built the tool. An institution occupying both positions carries both sets of obligations and must know which activity sits where.

> WHY IT IS A DISTINCT DISCIPLINE

The Builder can inspect its own system; the Buyer governs across an information boundary, with a contract, documentation of the vendor's choosing, and the system's observable behavior. Standard governance frameworks, written implicitly from the Builder's vantage, assume access the Buyer does not have. Buyer-side oversight is organized around what the Buyer can demand, observe, and prove.

PART TWO

The non-transfer principle, four times over

2.1 Banking: the mature statement

Banking answered the central Buyer question more than a decade before generative AI raised it: the institution that deploys a model owns the model risk, whoever built it. The principle just survived a complete rewrite of the guidance.

Why it matters. A principle that survives a fifteen-year, ground-up revision is the durable layer, not a rule of the old framework. SR 26-2 also handed every board a finding worth more than the guidance itself: the agencies deliberately placed generative and agentic AI outside the revised framework. The most mature model-governance instrument in American supervision declines to cover the AI arriving fastest. The gap is not closing on its own.

RULE

Interagency Guidance on Third-Party Relationships, 88 Fed. Reg. 37920 (June 6, 2023); SR Letter 26-2, Revised Guidance on Model Risk Management (Apr. 17, 2026), superseding SR 11-7 (2011) while keeping the model inventory expectation and the non-transfer principle, moving to a risk-based cadence, and scoping generative and agentic AI out.

2.2 Employment: the bluntest statement

The employment decision is the employer's decision, and the vendor's role in producing it does not displace the employer's responsibility. The Buyer who deploys a screening tool it cannot explain has purchased an exposure it cannot inspect.

Why it matters. The 2025 withdrawal of the EEOC's algorithmic technical assistance changed federal enforcement posture on disparate-impact theories; it changed nothing about the statutes, private rights of action, or the states, several of which now reach the Buyer directly.

AUTHORITY

Title VII, the ADA, and the ADEA; Illinois H.B. 3773 / Pub. Act 103-0804 (eff. Jan. 1, 2026), reaching unintentional disparate impact, barring ZIP-code proxies, and making the employer (not the vendor) answerable; NYC Local Law 144 (published bias audits since 2023); Connecticut Pub. Act 26-15, where the use of AI is no defense but documented anti-bias testing mitigates.

2.3 Healthcare: the identification duty

Federally funded health programs must make reasonable efforts to identify the patient-care decision-support tools they use, and mitigate the discrimination risk. The duty is written for Buyers in everything but name.

Why it matters. The tools it reaches, risk scores, sepsis predictors, scheduling optimizers, arrive embedded in the electronic health record and in vendor software the institution adopted without evaluating as AI. A health system that strips its tool inventory because a rule may be revised has confused the obligation with the instrument that described it; the discrimination was unlawful before the rule named the tools.

> **RULE**
>
> ACA Section 1557 final rule, 89 Fed. Reg. 37522 (May 6, 2024), decision-support-tool provisions; compliance date May 1, 2025. Enforcement posture is uncertain under administrative review, but Section 1557 is a statute.

2.4 Insurance: the program expectation

The insurer's written AI program is examined on the bought AI, not only the built. Insurance supervision states the non-transfer principle in its most operational form.

Why it matters. Over half the states now expect a written AI program covering third-party AI, with documentation producible in examination and an examiner tool arriving in late 2026. Colorado's life-insurance regime, binding the use of external data and models, survived the repeal of the state's comprehensive AI statute, evidence that sector supervision is the stable track.

> **AUTHORITY**
>
> NAIC Model Bulletin on the Use of AI Systems by Insurers (Dec. 4, 2023), adopted in 24 states and D.C.; Colorado Division of Insurance Regulation 10-1-1 (3 CCR 702-10) under SB 21-169.

Four regimes, four vocabularies, one principle: the deployment is the institution's act, and the accountability for it never signed the vendor contract.

QUESTIONS FOR THE BOARD

- For each AI system we did not build, who inside the institution is accountable for its outcomes, and can they suspend it?
- Does our vendor diligence read the safety frameworks and incident reports the Builder statutes now require developers to publish?

PART THREE

The Buyer's program

If accountability does not transfer, the Buyer needs a program built for its actual vantage point: a vendor-embedded inventory, named ownership, escalation through the vendor relationship, and a record kept with built-AI discipline.

Why it matters. Each element answers what must be governed and who is accountable, with implementation left to the operational layers. The elements are the Caremark line's anatomy, translated to the Buyer who governs across an information boundary.

> THE FOUR ELEMENTS
>
> **Vendor-embedded inventory:** the AI that entered through enterprise software acquired before anyone evaluated it as AI, captured with the system, the decisions it influences, the vendor, and the statutory role the institution occupies. **Named ownership:** an officer who knows what the system does, sees its outcomes, can suspend it, and carries the McDonald's upward duty, ownership that cannot live in procurement or with the vendor. **Escalation through the relationship:** contractual incident flow-through on defined clocks (72 hours in New York, fifteen days in California for frontier incidents; four business days for material cyber incidents), meeting the EU AI Act's deployer monitoring and log-retention duties. **The record:** inventory currency, diligence and its sources, the owner's monitoring, escalations and responses, designed rather than accumulated, and rewarded by the Texas and Connecticut documentation defenses, which do not distinguish built from bought.

CONCLUSION

Build the program as if the AI were yours

Buying is deploying, deploying is the institution's act, and the act requires the same program the Caremark line demands for any mission critical risk.

The Buyer position feels safer than it is. The institution that builds nothing assumes it has delegated its AI risk along with its AI, and four bodies of law say otherwise in unison. Accountability stays home. Build the program as if the AI were yours, because in every forum that matters, it is. The program is the governance.

Sources and Further Reading

Authority: Interagency Guidance on Third-Party Relationships, 88 Fed. Reg. 37920 (June 6, 2023); SR Letter 26-2 (Apr. 17, 2026); Title VII; ADA; ADEA; Illinois Pub. Act 103-0804; NYC Local Law 144; Connecticut Pub. Act 26-15; ACA Section 1557 rule, 89 Fed. Reg. 37522 (May 6, 2024); NAIC Model Bulletin (Dec. 4, 2023); Colo. Div. of Ins. Reg. 10-1-1; Regulation (EU) 2024/1689 (deployer obligations, Art. 26); Cal. Bus. & Prof. Code § 22757.10 et seq.; New York RAISE Act; Tex. Bus. & Com. Code chs. 551-552; In re McDonald's Corp. Stockholder Derivative Litig., 289 A.3d 343 (Del. Ch. 2023).

Doctrinal treatment of the oversight cases appears in the Center's *Caremark AI Liability Roadmap*; current status of every obligation cited is maintained in *The AI Oversight Obligations Reference*.