

The Missing Oversight Layer in the Trustworthy AI in Critical Infrastructure Profile

Comment on the NIST TACIP Community of Interest Discussion Draft of 3 August 2026 · Brian J. Allen, Esq., Executive Director · cfaio.org

BOTTOM LINE UP FRONT

TACIP is strong at the operational and technical layers and does not establish the layer above them.

The profile defers scope, granularity, and reporting output to the organization in place after place, correctly and by design. What it never says is who in the organization holds those decisions, or on what record. It never states who is accountable for an AI deployment, who grants the authority to operate it, or who can withdraw it. We propose a Practice 13 that establishes the AI oversight program by attaching to the designated-accountability mechanism each sector already has, rather than creating a new one.

Four findings you can check in ten minutes

- 1. Implementation 3.3.3 routes operators to "Escalate to Oversight." The profile never defines Oversight.** Not what it is, who staffs it, what authority it holds, or what it must do on receipt. The other three response pathways resolve to a defined action by a defined party. This one resolves to a word.
- 2. Practices 3 and 5 both carry the Govern 2.3 mapping, executive leadership responsibility for decisions about AI risks. No task under either assigns anything to executive leadership.** Across the table, "executive" appears in two substantive places: those two mapping labels, and Task 9.1's training range. Anyone crosswalking against the AI RMF will see it.
- 3. OCC Bulletin 2026-13, which the reviewer note at 2.3.4 is weighing, expressly excludes generative and agentic AI from its scope.** "Generative AI and agentic AI models are novel and rapidly evolving. As such, they are not within the scope of this guidance." Worth citing with that caveat.
- 4. ISO/IEC 38507:2022 appears nowhere in the draft, including the tracking list.** It is the one international standard written for members of a governing body on the organizational use of AI.

The finding

Read as a set, the profile's twelve Practices are twelve things an organization *does*. Each is a statement about activity. None is a statement about accountability. The unit of oversight governance is not a control, a committee, or a policy. It is a program: a documented structure stating what the organization is governing, what it is positioned to decide, and who is answerable. That program is distinct from the control inventory and the risk management process and does not emerge from either. Practices 1 through 12 are the machine. The oversight program is the specification for the machine, and a machine cannot specify itself. Absent that specification, the settings resolve to whoever is closest to the keyboard.

The practical test. An operator can implement every Task and Implementation in Practices 1 through 12 and still be unable to answer the four questions an examiner will ask: what does this AI system put at risk, who is accountable for it, under what authority is it operating, and what would cause you to withdraw that authority.

The proposal: Practice 13, drafted for insertion

Establish, charter, and evidence the AI oversight program. The charter states what the program governs, at what granularity the profile is applied, what authority the program carries, and what it must produce. Seven tasks run from it:

13.1–13.2

Charter and named accountability. One instrument, on a record the organization already maintains, and an accountable officer designated through the sector mechanism where one exists.

13.3–13.4

Tolerance and authority to operate. Thresholds ratified at the accountable level; authority to operate granted as a dated, scoped, time-bound instrument, with human-out-of-the-loop operation authorized explicitly.

13.5–13.7

Reporting, assurance, competency. A decision-ready report that reaches the officer; independent review outside the operating line; competency defined by the decisions the officer makes.

Minimum viable version. For a resource-constrained operator, Tasks 13.1, 13.2, and 13.5 carry most of the value: a one-page charter, a named officer, and a report that reaches that officer. An organization that does only those three has an oversight program; an organization that does none of them has controls without an owner.

Why the mechanism already exists

The comment maps the designated-accountability mechanism in all sixteen critical infrastructure sectors. Eight carry a mandatory federal requirement to designate a named accountable individual. Four carry a partial or role-based analog. One is entity-level only. Three have none. Fifteen of the sixteen instruments are silent on AI; the exception is Government Services and Facilities, where OMB M-25-21 requires each agency to designate a Chief AI Officer. The federal government has imposed on itself the named AI accountability requirement it has imposed on none of the sectors it regulates. Routing to the existing mechanisms is cheap. Waiting for sixteen regulators is not.

The precedent

The 2014 Cybersecurity Framework contained no GOVERN Function. The control layers matured first, the oversight layer went unbuilt, and the gap was closed a decade later from more than one direction: GOVERN added in CSF 2.0 (February 2024), SEC governance disclosure requirements (July 2023), and the *Caremark* oversight duty applied to mission-critical risk in *Marchand v. Barnhill* (Del. 2019). None of that displaced the control frameworks. It was added on top, after the fact, on terms the standards community did not set. AI compresses the interval. TACIP can define the oversight layer now, in its own vocabulary, or it can be defined later by someone else.

Also in the comment

Eighteen indexed comments with concrete alternative text keyed to the profile's numbering (priority rows: 3.3.3, 3.1.3, 4.5.2, 6.1.2, and the "Do no harm" criterion). Answers to the glossary, design criteria, and prioritization questions put to reviewers. A proposed locatable-owner design criterion. Reference notes, including the OCC scope caveat and the ISO/IEC 38507 addition. An offer to draft the missing Appendix C operational scenario for a resource-constrained operator.

The Center for AI Oversight is a 501(c)(3) nonprofit defining oversight governance for AI in regulated industries. It addresses two questions for boards, audit committees, and senior executives: what must be governed, and who is accountable. It does not prescribe how. The Center has no commercial interest in any vendor, product, or control framework referenced in this comment. The full 20-page comment includes drafted Practice 13 text in the profile's format, the sixteen-sector appendix, and the disclosure of AI assistance. Contact: Brian J. Allen, Esq., Executive Director · cfaio.org