

COMMENT TO THE COMMUNITY OF INTEREST

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

A proposed Practice 13, with indexed comments on the Discussion Draft of 3 August 2026

SUBMITTED TO	Raymond Sheh and Martin Stanley, National Institute of Standards and Technology aiciprof@nist.gov
SUBJECT	NIST AI RMF Profile on Trustworthy AI in Critical Infrastructure (TACIP), Community of Interest Discussion Draft, 3 August 2026
SUBMITTED BY	Brian J. Allen, Esq., Executive Director, Center for AI Oversight Represented the banking sector, through the Bank Policy Institute, in the development of the NIST AI Risk Management Framework. Appointed by the Chairman of the Federal Communications Commission to represent the communications sector during the development of the NIST Cybersecurity Framework.
DATE	11 August 2026
RESPONDS TO	Strategic Governance and Executive Translation (Open Design Area); Practice 13 solicitation; Glossary Priorities; Document Design Criteria; Ideas and References Already in Progress

The Center for AI Oversight is a 501(c)(3) nonprofit defining oversight governance for AI in regulated industries. This comment is offered as public input to a discussion draft. It does not represent official guidance of NIST or of any agency.

BOTTOM LINE UP FRONT

TACIP is strong at the operational and technical layers and does not establish the layer above them. What the profile does not set out is the program: what the organization is governing, what it is positioned to decide, and who is answerable. That program is distinct from the control inventory and from the risk management process, and does not emerge from either.

The profile defers the scope, the granularity, and the reporting output to the organization in place after place, correctly and by design. What the profile never says is who in the organization holds those decisions, or on what record. It never states who is accountable for an AI deployment, what that deployment puts at risk, who grants the authority to operate it, or who can withdraw it.

We propose a Practice 13 that establishes that program by attaching to the designated-accountability mechanism each sector already has, rather than creating a new one. Eight of the sixteen sectors in Table 1 already carry a mandatory federal requirement to designate a named accountable individual, and four more carry a partial or role-based analog. Fifteen of the sixteen instruments are silent on AI. The appendix maps all sixteen. The profile should route to those mechanisms rather than assume the reader will supply them.

This is the Center's first submission to the TACIP community of interest, keyed to the Community of Interest Discussion Draft of 3 August 2026. It is organized in five substantive parts. Part 1 states the finding and answers the question at the Practice 13 placeholder about why a proposed risk area needs its own Practice. Part 2 supplies drafted Practice 13 text in the profile's Practice, Task, and Implementation format. Part 3 provides eighteen indexed comments keyed to the profile's own numbering. Part 4 answers the direct questions put to reviewers on glossary priorities, design criteria, and roadmap prioritization. Part 5 annotates and adds references, including a scope caveat on a reference the current text is weighing.

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. It is a routing label with no destination, and it is the profile’s own word for the layer this comment proposes. Part 3, Implementation 3.3.3 row.

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 rows of the table of Practices, Tasks, and Implementations, “executive” appears in exactly two substantive places: those two mapping labels, and Task 9.1’s reference to training “from executive level to front-line operations.” Every other occurrence is a reference title. Anyone crosswalking the profile against the AI RMF will see it. Part 3, Practice 3 and Practice 5 rows.

3. OCC Bulletin 2026-13, which the reviewer note at 2.3.4 is weighing, expressly excludes generative and agentic AI from its scope. Part 5.1.

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. Part 5.2.

A two-page summary sheet accompanies this comment and is intended to be circulated separately from it.

The Center’s Executive Director is available for a follow-up conversation with the team, as invited in the Tips for Providing Feedback.

1 The finding, and why the oversight layer needs its own Practice

Section 1.1 of the draft states that Practices are written for strategic or executive-level managers, Tasks for technical managers, and Implementations for the personnel who carry out the task. That is a hierarchy of altitude within one discipline. It is not a hierarchy of layers. Read as a set, the twelve Practices are twelve things an organization *does*: establish requirements, define expectations, define policies and procedures, maintain override and recovery procedures, implement identity controls, integrate supply chain visibility, manage provenance, incorporate incident response, provide training, implement logging, maintain continuity, validate artifacts. Each is a statement about activity. None is a statement about accountability.

This is a gap rather than an error, and the recommendation is additive. The management layer in the current draft is well built and should stay. What the profile needs is an oversight layer in addition to it, kept distinct rather than written as a higher-altitude version of it. The distinction is easy to lose, because a management statement written abstractly enough begins to read like a governance statement. Altitude is not accountability. A Practice can be pitched at the C-suite and still describe only what the organization does.

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 for those decisions. That program has to stand on its own, distinct from the control inventory and from the risk management process. It does not emerge from either, however completely they are built. An organization can hold a full inventory of AI controls and a mature risk management process and still be unable to demonstrate that it is positioned to make an informed decision about the risk and reward of a given deployment.

The draft acknowledges the boundary. The Introduction notes that the profile reaches “upward towards governance to a shared practice layer” while remaining “intentionally not ‘governance-complete.’” We are not asking the profile to become governance-complete, or to restate ISO/IEC 38507 or ISO/IEC 42001. We are asking for the narrow band of connective tissue that lets an accountable officer use the twelve Practices as instruments of accountability rather than as a description of good engineering.

1.1 What the oversight program does that Practices 1 through 12 cannot

One point before the three functions, because it determines how the Practice must be written. The duty to own an AI oversight program is not created by regulation. Where a governing body exists, it arises from that body's own duty to oversee mission-critical risk. Where a license or statute applies, regulation adds a second and more specific basis. Where neither reaches, it remains the standard of care a board or an accountable executive is expected to meet. Regulation supplies the mechanism, not the reason.

That distinction matters practically. Three of the sixteen sectors mapped in the Appendix have no federal designation mechanism at all, and an AI deployment can be safety-relevant or continuity-critical without touching any specific regulatory obligation. A Practice scoped only to regulated activity would exclude both. The Practice below is therefore scoped by consequence, with regulatory obligation as one trigger among several rather than the trigger.

Three functions, and each is an input to the twelve Practices rather than an output of them.

Scope. Section 1.2 applies the profile to all AI systems used in critical infrastructure environments, including general-purpose productivity and analytics tools adopted without CI-specific vetting. A universal scope statement is the correct starting point for a cross-sector profile, and it is not a scoping decision. Someone has to determine which deployments sit inside the organization's oversight perimeter, on the basis of what each puts at risk: safety, service continuity, the public, and where they apply, statutory and license obligations. That determination precedes Practice 1 and is not assigned anywhere in the draft.

Granularity. The profile defers the level of rigor to the organization, correctly and repeatedly. Section 1.1 has organizations “dynamically establishing which controls must be implemented immediately versus which can be phased in” and allocating “limited resources within their unique risk profile.” Section 1.2 states that organizations “may apply this profile at different levels of rigor.” Section 2 states that prioritization “is specific to the implementing sector or organization.” Implementation 1.3.1 sets TEVV scope “proportionate to consequence severity.” Task 5.1 states that the appropriate level of identity granularity “will vary depending on” available resources, expected threat level, expected impact, other mitigating controls, scale, vendor co-operation, and application complexity. Implementation 8.2.1 allows root cause analysis requirements to be relaxed with compensating controls. The Document Design Criteria ask that the reader be given “space to make this trade-off.” In each case the profile is right to leave the setting to the organization. In none of them does it say who in the organization holds the setting, or on what record.

Output. Practices 8 and 10 produce evidence and Task 10.3 reviews it. No practice specifies what has to reach a decision-maker, in what form, and on what cadence, in order to support a decision to continue, constrain, or withdraw a deployment. Evidence production and decision support are different requirements, and the profile currently supplies only the first.

THE OBJECTION WE EXPECT, AND THE ANSWER

TACIP already addresses governance, at the Practice layer, and risk management, in Practices 1 through 3. Both are true, and neither closes this gap. The Practice layer is where the profile makes the deferrals set out above. Practices 1 through 3 are where the risk management process runs against settings those deferrals hand to the organization. Practice 13 is not another instance of either. It is the row that says to whom the deferrals are addressed.

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.

1.2 Where the gap shows in the current text

Section 1.1 of the draft states that “Organizations superimpose their own top-level enterprise risk governance onto these concrete tasks.” That verb assumes the oversight layer exists and is already competent on AI. In the regulated critical infrastructure organizations the Center works with it usually is not. The governing structure predates AI adoption, assigns accountability for cyber and for safety, and makes no AI-specific assignment at all.

The front matter sharpens the point rather than closing it. Section 1.1 directs that “organizations, and executive leadership in particular, may consult Appendix B: Historic Incident-to-Practice Mapping ... before proceeding to the corresponding practice descriptions *to establish governance ownership* and delegate Task-level implementation.” The instruction is right. What follows it is a table in which nothing tells executive leadership how to establish governance ownership, what that ownership covers, or that it should be recorded anywhere. The profile asks for the artifact in its front matter and supplies no row that produces it.

The consequence appears throughout the table as decisions with no decision-maker. Implementation 3.1.3 directs the organization to establish the threshold criteria that bound acceptable AI behavior, and states it without an owner. The text directs that setting away from CVSS and confidence scores toward consequence-driven prioritization against “real-world impacts to public health, safety, and mission continuity,” which is right and raises rather than lowers the need for a ratifying party. Task 12.1 and Implementation 12.1.1 both route to 3.1.3 for consequence tiering; elsewhere the profile sets thresholds locally, at Implementation 3.3.2 for drift and Implementation 8.2.2 for safety margins, none of them ratified anywhere either. Implementation 6.1.2 directs that unresolvable supply chain documentation gaps be “escalated and factored into the organization’s go/no-go deployment decision,” without naming the escalation destination or whose decision the go/no-go is. Implementation 1.3.1 constrains a system to increased human oversight “until TEVV evidence supports a higher level of autonomy,” without identifying who judges that it does. Implementation 8.3.2 forms a permanent AI incident governance team of the Chief Information Security Officer, technical staff, legal staff, and operations management, with no line to an accountable officer or a governing body.

The clearest instance is Implementation 3.3.3, which gives operators four pre-authorized response pathways, one of which is “Escalate to Oversight,” for persistent anomalies that indicate drift or data schema shifts. Nothing in the profile says what Oversight is. The other three buckets resolve to a defined action by a defined party. This one resolves to a word. Practice 13 supplies its destination, and if the Practice is not adopted, 3.3.3 still needs one.

The contrast with Practice 5 is instructive. Implementation 5.1.1 requires cryptographically verifiable, non-shared identity for every autonomous sub-agent and transient tool invocation, on the ground that shared credentials obscure the root cause of an action and defeat granular attribution. The reasoning is correct. Applied consistently, it also reaches the humans who authorize those agents to act on regulated processes.

As drafted, the profile holds AI agents to a higher standard of identifiable attribution than the people accountable for deploying them. Practice 5 also carries the Govern 2.3 mapping, and none of Tasks 5.1 through 5.4 assigns anything to executive leadership. The pattern at Practice 3 is therefore not an isolated drafting oversight in one row. It recurs wherever the mapping is applied.

THE PRACTICAL TEST

A regulated critical infrastructure 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, and which obligations or duties does it touch, who is accountable for it, under what authority is it operating, and what would cause you to withdraw that authority. Practice 13 exists to make those four questions answerable.

1.3 Why this requires its own Practice

The Practice 13 placeholder asks contributors to explain why a proposed risk area needs its own Practice rather than sitting under an existing one. Four reasons.

The audience is different, and the design criteria say so. The criteria specify that a Practice should be readable by a non-technical manager who will then direct the work, and that each Practice should let the reader identify which internal team leads the objective. This material is not directed at a team. It is directed at the officer who is answerable when the team's work is examined. Placed as a Task under Practice 3, it would be read by a technical manager with no authority to assign accountability upward.

It answers a different question. Practices 1 through 12 each state something the organization does about AI risk, which is why the draft is right that their order is arbitrary and that any one is a valid entry point. This Practice states who is answerable for those activities and on what record. It is a valid entry point on the same terms, and no reordering of the others is implied. Distributing it as a sentence appended to several existing rows would reproduce the current situation, in which accountability is implied everywhere and required nowhere.

It is the most transferable content in the profile, not the least. Operational realities diverge sharply across the sixteen sectors: a water utility, a clearing bank, and a nuclear licensee share almost no operational vocabulary. Their accountability structures do not diverge in the same way. The Appendix maps the designation mechanism in each of the sixteen sectors. Eight carry a mandatory federal requirement to designate a named individual at the program or officer level: Communications, Dams, Defense Industrial Base, Energy, Financial Services, Government Services and Facilities, Healthcare and Public Health, and Transportation Systems. Four more carry a partial or role-based analog. One is entity-level only. Three have no federal designation mechanism at all. The mechanism recurs in three quarters of the sectors and fails in a quarter, and the profile should say which is which. A cross-sector profile is the correct place to state the mechanism and to identify where it is absent. A sector companion is the correct place to state which office fills it.

Sector supervision is moving away from supplying this layer, not toward it. Financial services has the deepest model risk governance tradition of any critical infrastructure sector. On 17 April 2026 the Federal Reserve issued SR 26-2, the Office of the Comptroller of the Currency issued Bulletin 2026-13, and the Federal Deposit Insurance Corporation issued FIL-15-2026, together constituting the revised interagency guidance on model risk management. That guidance expressly excludes the systems this profile principally addresses: "Generative AI and agentic AI models are novel and rapidly evolving. As such, they are not within the scope of this guidance." The most mature sector-level instrument in critical infrastructure has just

carved out the technology. A cross-sector profile that assumes sector regulation will supply the oversight layer is assuming something currently untrue in the sector where it was most nearly true.

1.4 The sequencing risk

You know this history better than we do, so we will state it once. The 2014 Cybersecurity Framework organized its Core around five Functions and contained no GOVERN Function. The control layers matured first, the oversight layer went largely unbuilt, and the gap was closed a decade later from more than one direction: GOVERN was added as a sixth Function in CSF 2.0 in February 2024, the SEC adopted governance disclosure requirements in July 2023 including board oversight at Item 106(c)(1) of Regulation S-K, and in the courts the *Caremark* oversight duty was applied to mission-critical risk in *Marchand v. Barnhill*, 212 A.3d 805 (Del. 2019), a mechanism that reaches only operators with a corporate board and many in these sectors have none.

What is less documented is what it produced at the receiving end, and that is offered here as field observation. In assessing and building cybersecurity oversight practices for organizations globally at EY, and in later work with operators, executives, and boards across sectors, the pattern does not vary by sector or by size. Asked what their own oversight role is, they point to the Cybersecurity Framework. It is the artifact everyone reaches for, and it does not answer that question, because it was not written to. It describes what the organization does about the risk, not what the accountable level is answerable for governing, at what granularity, or what the program owes them. There has been nothing to point to, and that absence is the argument for populating this layer now rather than later.

None of that displaced the control frameworks. It added a layer above them, after the fact, on terms the standards community did not set. AI compresses the interval, because the standard of care is arriving ahead of the controls rather than a decade behind them. TACIP can define the oversight layer now, in its own vocabulary and with its community of interest, or it can be defined later by someone else.

2 Proposed Practice 13, drafted for insertion

The following is offered as drop-in text in the profile's format. Cross-references use the draft's existing index numbers. Implementations are paired to give a minimal and a fuller option, per the Document Design Criteria. Where the profile's convention is to route rather than restate, this text routes.

PROPOSED PRACTICE 13: Establish, charter, and evidence the AI oversight program

Charter the AI oversight program, then operate it. The charter states what the program governs, at what granularity this profile is applied, what authority the program carries, and what it must produce. From the charter run the named accountability for each affected obligation, the risk tolerance within which deployments may operate, the time-bound authority to operate, the reporting and disclosure line, independent assurance over the whole, and the competency required of the officer who signs.

Organizations should attach this practice to the designated-accountability mechanism their sector already uses rather than create a parallel AI governance structure. A structure that is not examined will not be maintained. The transferable element across sectors is the mechanism, not its content.

Routing. This practice is led by the function that maintains the organization's governance instruments, which is typically the corporate secretary, the general counsel, or the office that maintains the safety or security program charter. The risk function supplies the tier criteria at 13.1.2,

the function operating the AI system supplies the register content at 13.1.3, and internal audit or its equivalent performs Task 13.6. No new department is required.

Minimum viable. For an operator with limited resources, 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. Tasks 13.3, 13.4, 13.6, and 13.7 are the fuller form and can be phased in. An organization that does only the first three has an oversight program; an organization that does none of them has controls without an owner.

This practice does not duplicate or replace the operational and management controls in Practices 1 through 12, and requires no restructuring of them. It sits above them and states the conditions under which they are performed by a party who is answerable for them. The indexed comments in Part 3 propose targeted text changes to several existing rows; those are separate suggestions, each of which stands on its own, and none is a precondition for adopting this Practice.

Related Keywords: Charter and mandate; scope boundary; granularity tiers; decision rights; accountable officer; designated accountability; delegation of authority; authority to operate; risk appetite and tolerance; periodic attestation; escalation and disclosure; independent assurance; residual risk acceptance.

Main Trustworthiness Characteristics and Terms: Accountable & Transparent (primary); Safe; Resilient.

Primary AI RMF Mappings: Govern 1.1 (legal and regulatory requirements); Govern 1.3 (determining the needed level of risk management activity based on organizational risk tolerance); Govern 2.1 (roles, responsibilities, and communication); Govern 2.3 (responsibility of executive leadership for decisions about AI risks); Govern 3.2 (roles and responsibilities for human-AI configurations and oversight); Govern 4.1 (critical thinking and safety-first mindset); Govern 1.4 (risk management process established through transparent policies, procedures, and controls, which is the charter); Govern 1.5 (ongoing monitoring and periodic review planned, with roles defined and review frequency determined); Govern 1.6 (mechanisms to inventory AI systems, resourced according to risk priorities); Govern 2.2 (personnel receive AI risk management training, which is Task 13.7); Map 1.5 (organizational risk tolerances); Manage 1.4 (negative residual risks documented to downstream acquirers and end users).

Primary Cyber AI Profile Mappings: To be determined on publication of NIST IR 8596. Note that the Cyber AI Profile is expected to address the security of AI systems rather than the accountability of the organizations operating them, so this Practice is unlikely to be superseded by it.

13.1 Charter the AI oversight program: its mandate, its scope, its granularity, and its required outputs

Prepare, and maintain for approval at the accountable level, a single instrument stating what the program governs, on what basis, what authority it carries, at what granularity this profile is applied, and what the program must produce. Every other task in this Practice runs from settings this task records. This is the constituting act; the remaining tasks are its sections rather than its siblings.

13.1.1 Charter on an instrument the organization already maintains (minimal)

Amend the charter, terms of reference, or designation record the organization already keeps to state the program's purpose, the classes of AI activity in scope, the authority to require information and to pause or withdraw a deployment, the approving party, and a review date. One page attached to an existing document. This is deliberately not the formation of a new committee: a committee is evidence of attention, not of accountability.

State the basis on which the program is chartered. In most organizations more than one applies: the governing body's own duty to oversee mission-critical risk; a statutory, license, or tariff obligation where one reaches the activity; and the standard of care expected of an operator of

critical infrastructure. Naming the basis matters because it determines what happens to deployments that touch no specific regulation, which is most of them in several sectors.

Where the sector supplies a governing or licensing authority, name it as the chartering party. Where the organization is a municipal utility, a cooperative, or a district with no corporate board, the chartering party is the body that already approves its safety or security program. Where the sector supplies no designation mechanism at all, as in three of the sixteen, the governing body charters directly on its own authority. The absence of a regulatory mechanism removes the plumbing, not the duty.

13.1.2 Set the scope boundary and the granularity tiers, and record who set them (minimal)

This is the task that answers the profile's own deferrals. Define the criteria that place a deployment inside the oversight perimeter. Those criteria should turn on consequence, meaning what the deployment puts at risk in safety, service continuity, and public impact, with statutory and license obligations as an additional trigger rather than the only one. Define the consequence tiers that determine the level of rigor applied. State, for each tier, which of Practices 1 through 12 apply in full, which apply in their minimal implementation, and which are deferred with compensating controls. Name the officer who set the tiers and record the basis.

This is where the profile's deferrals land. Section 1.1, Section 1.2, Section 2, Implementation 1.3.1, Task 5.1, and Implementation 8.2.1 each defer a level of rigor to the organization, correctly. Part 1.1 above sets them out. None says who holds the setting.

13.1.3 Record what is at stake and who is accountable, against the inventory (minimal)

Extend the Master AI Asset List maintained under Implementation 7.1.1 with two fields: what the deployment puts at risk, and the accountable officer. The first field may hold a statutory or license obligation, a safety case, a service continuity dependency, or a duty owed by the governing body. Where a compliance obligation register exists, this is a join between two records the organization already keeps. Where it does not, the field is populated from the safety and continuity analysis the organization performs anyway.

Re-run on any new deployment, any material model, vendor, or data change under Task 7.5, and any change in the underlying obligation. Section 1.2 brings general-purpose productivity and analytics tools into scope, and this is the control that determines which of those entered the oversight perimeter by accident.

13.1.4 Declare the program's required outputs (fuller)

Enumerate what the program must produce and maintain: the obligation and officer register under 13.1.3; the designation and delegation record under 13.2; the ratified tolerance and its basis under 13.3; the authorization register with expiries under 13.4; the oversight report under 13.5; the residual risk acceptance log under 13.2.2; and the independent assurance result under 13.6.

An output that no instrument requires will not be produced under resource pressure, and its absence will not be noticed until examination.

13.2 Assign and record named accountability

Using the scope and authority recorded in the charter under Task 13.1, prepare, for designation at the accountable level, the officer accountable for each deployment inside the perimeter, using the designation mechanism the sector already provides where one exists and the governing body's own authority where it does not. Record delegations, their scope, and their currency. This task answers the question of who is accountable.

13.2.1 Attach to the existing sector mechanism (minimal)

Where the sector provides a designation mechanism, extend its scope to AI-affected obligations rather than create a new AI accountability structure. Examples, given for structure rather than content: NERC Reliability Standard CIP-003-9 Requirement R3 requires each Responsible Entity to identify a CIP Senior Manager by name and to document any change within thirty calendar days, with delegations documented, approved, and updated under Requirement R4. The HIPAA Security Rule at 45 C.F.R. 164.308(a)(2) requires a covered entity or business associate to identify the security official responsible for the development and implementation of the policies and procedures required by the subpart. TSA Security Directive Pipeline-2021-01G, Section II.B.1.a, requires owner/operators to designate and use a primary and at least one alternate Cybersecurity Coordinator at the corporate level, each required under Section II.B.2.b to be accessible to TSA and CISA twenty-four hours a day, seven days a week.

The transferable pattern is a named individual, a documented and time-bounded delegation path, and a currency requirement. Scale the apparatus to the organization. A large operator may maintain a delegation register with defined refresh cycles; a small municipal utility, cooperative, or district may satisfy this with a single named individual recorded on the instrument it already uses for its safety or security program and reviewed when that instrument is reviewed. What does not scale down is the naming itself, which costs a line of text and is what the remainder of this Practice depends on.

13.2.2 Record residual risk acceptance, including out-of-scope declarations (fuller)

Where a control in this profile is not implemented, is partially implemented, or is declared inapplicable under the "Do no harm" design criterion, record the acceptance with the name and role of the accepting officer, the scope of the acceptance, the compensating controls relied upon, and a review or expiry date. An out-of-scope declaration with no named acceptor and no expiry is indistinguishable from an undocumented gap. Where the sector supervises the activity it will be read as one. Where it does not, the organization has no way to establish later whether the exclusion was a judgment someone made or an omission nobody noticed.

13.3 Set and ratify AI risk tolerance at the accountable level

Risk tolerance is an oversight decision. The thresholds established under Implementation 3.1.3, the evaluation cadence set under Implementation 2.3.4, the drift boundaries set under Implementation 3.3.2, and the restoration priorities set under Task 11.1 all resolve against a tolerance that the profile currently leaves unowned.

13.3.1 Ratify thresholds and record the ratification

Ratify the thresholds established under Implementation 3.1.3 at the level of the officer designated under Task 13.2, and record the date, the basis, and the evidence relied upon. Ratification is

distinct from review: review confirms a threshold is still described accurately, ratification accepts the consequences of operating at it.

13.3.2 Define re-ratification triggers

Define the conditions that require re-ratification rather than review: a material change to the deployment or its operating environment, an incident meeting the criteria identified under Task 8.1, a drift excursion beyond the thresholds set in Implementation 3.3.2, a change in the underlying regulatory obligation, or a defined elapsed period aligned to the organization's existing risk governance cycle.

13.4 Grant, bound, and periodically re-ratify authority to operate

Introducing AI into a regulated critical infrastructure process is an authorization decision with a beginning, a scope, and an end. The profile describes several conditions under which autonomy is constrained or expanded, at Implementations 1.3.1, 4.5.1, and 12.1.3, but does not identify who grants that authority, on what record, or how it is withdrawn.

13.4.1 Prepare the authority to operate as a dated, scoped, and time-bound instrument, for issue at the accountable level

Record the deployment, the process it affects, the autonomy regime authorized, the conditions of operation, the compensating controls relied upon, the accountable officer, and the expiry. Expiry matters more than approval: an authorization that does not lapse cannot be allowed to lapse, and becomes a permanent grant made once on evidence that has since aged.

13.4.2 Treat operation without real-time human review as a separate and more consequential grant

Task 4.5 correctly identifies that human oversight becomes nominal where the required review latency exceeds the system's action time. Implementation 4.5.1 documents that condition and Implementation 4.5.2 substitutes deterministic boundary constraints for human intervention. That substitution is presently framed as an engineering conclusion. It is an authorization decision, and it is the point at which the profile's human oversight language ceases to be load-bearing.

Where an organization operates in a human-out-of-the-loop regime, the regime should be authorized explicitly by the officer designated under Task 13.2, in a dated and time-bound instrument, and re-ratified on a defined cadence. Documenting that human intervention cannot be relied upon identifies a condition. It does not authorize operating under it.

13.4.3 Define automatic suspension conditions

Define, for each authorization, the conditions under which it lapses without further decision: expiry; a drift excursion beyond Implementation 3.3.2 thresholds; an incident meeting the criteria identified under Task 8.1; or the loss of a compensating control on which the authorization was conditioned.

What lapsing does to the process should be set by the process and not by this task. Where the fallback state defined in Implementation 3.2.3 is validated as safe to enter from the process state obtaining at the trigger, lapse returns the process to it. Where entering that state would itself be hazardous at that moment, or where the process has no safe intermediate state, the authorization should convert instead to a bounded continuation under named interim authority,

with a maximum duration and a scheduled transition to the alternative operating mode under Task 11.2. Automatic termination is appropriate for a decision right. It is not appropriate for a physical process, and this implementation should not be read as requiring an uncommanded transition.

13.5 Establish the oversight reporting, escalation, and disclosure line and its cadence

Practices 8 and 10 and Implementation 2.3.4 generate audit-ready evidence. The profile does not require that this evidence reach anyone with authority to act on it. Evidence that terminates in the technical management function is monitoring. It becomes oversight only when it reaches a party who can withdraw the authorization.

13.5.1 Route defined evidence categories to the accountable level

Route to the officer designated under Task 13.2, and where the organization has a governing body to that body, on a stated cadence: authorization status and pending expirations under Task 13.4; threshold excursions under Implementation 3.3.2 and sustained or recurring anomalies escalated under Implementation 3.4.4; incidents and near misses under Practice 8; residual risk acceptances outstanding under Implementation 13.2.2; and inventory coverage under Task 13.1.

13.5.2 Apply a decision-readiness test to the report

The report should support a decision to continue, constrain, or withdraw an authorization. A report that describes activity without supporting a decision does not satisfy this task, however complete its telemetry.

13.6 Establish independent assurance over the AI oversight structure

Practices 1 through 12 are performed by the functions that build, operate, and manage AI systems. No practice in the profile provides for a review independent of those functions. The Document Design Criteria state that outcomes should be verifiable, measurable, and audit-ready, and should enable regulatory oversight and forensic investigation. Audit-readiness without an independent reviewer is a property asserted rather than tested.

13.6.1 Bring the oversight structure into an independent review

Include the AI oversight structure in the scope of the organization's internal audit plan or, where no internal audit function exists, in a defined independent review by a party outside the operating line. Minimum scope: currency of the Task 13.1 charter and the scope and granularity settings it records, completeness of the 13.1.3 obligation register, currency of Task 13.2 designations, ratification of Task 13.3 thresholds, validity of outstanding Task 13.4 authorizations, and operation of the Task 13.5 reporting line.

13.6.2 Place assurance outside the operating line

Where the organization uses the Three Lines Model, place AI oversight assurance in the third line. First-line self-assessment against this profile is a useful management control and is not assurance.

In resource-constrained organizations the same small team may perform all three roles, and a mandated external review would be neither affordable nor proportionate. A gradient applies, in

descending order of independence: an external reviewer; a reviewer from elsewhere in the organization with no role in the AI deployment; a member of the governing or approving body outside management; or a peer review arranged through a sector ISAC or mutual-aid agreement, which Task 8.5 already identifies as an interim approach where specialist capability is scarce. Record which of these was used and why the next step up was not available, so that the degree of independence obtained is visible rather than assumed.

13.7 Establish competency at the accountable level

Practice 9 develops competency for operators, technical managers, and personnel who supervise AI systems. It does not develop competency for the officer who signs the authorization. This is a distinct requirement, and a narrower one: the officer needs enough to interrogate a report and test an assertion, not enough to build or validate a model.

13.7.1 Define competency by the decisions the officer makes

Define the minimum competency for the officer designated under Task 13.2 in terms of the specific decisions that officer makes: reading a TEVV result produced under Implementation 1.3.4, interpreting a drift excursion under Implementation 3.3.2, evaluating a residual risk acceptance under Implementation 13.2.2, and judging whether a report satisfies the decision-readiness test at Implementation 13.5.2. This fills the executive-level entry that Task 9.1 contemplates and Implementations 9.1.1 and 9.1.2 currently leave open.

13.7.2 Evidence competency at the accountable level

Record this competency in the same manner as Task 9.3, so that workforce readiness evidence covers the accountability layer and not only the operating layer. In several critical infrastructure sectors, competence at the accountable level is itself a supervisory expectation, and evidence of it is requested on examination.

3 Indexed comments on the current text

The following comments are keyed to the profile's own index numbers, per the request in the Community of Interest announcement. Suggested text is offered for adoption, adaptation, or rejection. Several rows reference Practice 13; each also identifies the defect in the current text independently, so the observations hold whether or not the Practice is adopted.

Index	Observation	Suggested change
Section 1.1 Audience	"Organizations superimpose their own top-level enterprise risk governance onto these concrete tasks." The verb assumes an oversight layer that exists and is competent on AI. In most regulated critical infrastructure organizations the governing structure predates AI adoption and assigns no AI-specific accountability.	Replace with: "Organizations apply their own top-level enterprise risk governance to these tasks. Where that governance does not yet assign accountability for AI specifically, Practice 13 sets out the structure to establish before the tasks below can be meaningfully assigned or resourced."

Index	Observation	Suggested change
Table 2 Trustworthiness	“Explainable & Interpretable” and “Accountable & Transparent” share a single definition, and that definition describes properties of the system: forensic telemetry, situational awareness, community trust. Accountability in the governance sense is a property of a person. A reader will reasonably conclude that accountability is discharged by logging.	Separate “Accountable & Transparent” into its own row: “Ensures that a named party within the organization is answerable for the introduction, operation, and withdrawal of the AI system, and that the supporting record is complete enough to be examined by a party outside the operating line.”
Practice 3 title and mappings	Practice 3 is titled to include oversight and maps to Govern 2.3, responsibility of executive leadership for decisions about AI risks. None of Tasks 3.1 through 3.8, including the newly added 3.8, assigns any responsibility to executive leadership. The mapping is asserted rather than operationalized, which will be visible to a reader crosswalking the profile against the AI RMF.	Move the Govern 2.1 and Govern 2.3 mappings to Practice 13, where a task discharges them, and retitle Practice 3 to “Define risks, policies, and operational supervision for automated AI and agentic behavior.” A task added under Practice 3 to discharge Govern 2.3 would be addressed to a technical manager who cannot assign accountability upward.
Practice 5 mappings	Practice 5 also carries the Govern 2.3 mapping, and none of Tasks 5.1 through 5.4 assigns anything to executive leadership. Across the entire table, “executive” appears in two substantive places only: the Govern 2.3 labels at Practices 3 and 5, and Task 9.1’s reference to training “from executive level to front-line operations.” The mapping is being used to signal that a Practice matters at the top of the organization rather than to assign anything there.	Apply the same treatment as at Practice 3: carry Govern 2.3 at Practice 13 and retain at Practices 3 and 5 only the mappings their tasks discharge. Where identity governance for AI entities is meant to reach the accountable level, the reachable statement is that the inventory of AI entities and their granted authorities is reported under Task 13.5, not that executive leadership is responsible for IdAM.
3.1.3	Risk thresholds are established with no owner named. The text correctly directs the setting away from CVSS and confidence scores toward consequence-driven prioritization against public health, safety, and mission continuity, which makes the missing ratification more consequential rather than less. Task 12.1 and Implementation 12.1.1 route to this row for consequence tiering.	Add: “Risk thresholds for AI systems affecting safety-critical or regulated critical infrastructure processes should be ratified at the level of the officer holding the organization’s designated accountability for the affected obligation, with the ratification recorded together with its date and basis. See Task 13.3.”
3.3.3	The four response pathways are a strong addition. Three of them resolve to a defined action by a defined party. “Escalate to Oversight” resolves to a word the profile does not define: not what Oversight is, who staffs it, what authority it holds, or what it must do on receipt. An operator following the quick-reference guide reaches a routing label with no destination, and the anomalies routed there are precisely the persistent, non-critical ones most likely to accumulate unaddressed.	Add to the “Escalate to Oversight” bucket: “Escalation under this pathway should be directed to the party identified in the AI oversight program charter as accountable for the affected deployment, on the reporting line and cadence established under Task 13.5, with a recorded disposition of continue, constrain, or withdraw. Where no such party has been identified, escalate to the officer holding the organization’s designated accountability for the affected regulatory or safety obligation.” The second sentence makes the row usable before Practice 13 is adopted.

Index	Observation	Suggested change
4.5.2	The substitution of deterministic boundary constraints for human intervention is presented as an engineering consequence of the time-constant finding at 4.5.1. It is an authorization decision, and the most consequential one in the profile.	Add: "Operation in a human-out-of-the-loop regime should be authorized explicitly, in a dated and time-bound instrument, by the officer accountable for the affected obligation, and re-ratified on a defined cadence. Documenting that human intervention cannot be relied upon identifies a condition; it does not authorize operating under it." See Implementation 13.4.2.
1.3.1	"Until TEVV evidence supports a higher level of autonomy" leaves open who determines that the evidence supports it, and on what record.	Add: "The determination that TEVV evidence supports a higher level of autonomy is an authorization decision and should be recorded as such, naming the deciding officer and the evidence relied upon."
6.1.2	Unresolvable supply chain documentation gaps are to be "escalated and factored into the organization's go/no-go deployment decision." Neither the escalation destination nor the owner of the go/no-go decision is stated. This is a second unrouted escalation, and unlike Implementation 3.3.3 it sits inside a deployment decision.	Replace the clause with: "escalated to the officer accountable for the obligation the deployment affects, and, where the deployment proceeds, recorded as a residual risk acceptance naming that officer, the compensating controls relied upon, and a review date."
7.1.1	The Master AI Asset List records "responsible contacts." A responsible contact is an operational point of contact, not an accountable officer, and the list carries no field for the regulatory obligations a deployment affects. An inventory built this way cannot support the go/no-go and authorization decisions the profile requires elsewhere.	Add two fields to the enumerated list: the accountable officer, and the regulatory obligation or obligations affected. See Implementation 13.1.3.
8.3.2	The permanent AI incident governance team comprises management, technical, legal, and operations functions. There is no escalation to an accountable officer or governing body, and no link to the external reporting obligations contemplated by Task 8.4.	Add: "and defines the threshold at which the officer accountable for the affected obligation and, where applicable, the governing body are notified, together with the external reporting obligations triggered under Task 8.4."
8.4.1 8.4.2 both TBD	These are the profile's external regulatory disclosure rows and are the rows most tightly coupled to named accountability, because in most sectors the person who signs a regulatory notification is designated by regulation and the reporting clock starts on a defined event.	Suggested content: for each deployment inside the oversight perimeter under Task 13.1, identify the reporting obligations a failure would trigger, the event that starts the reporting clock, the clock duration, and the officer authorized to sign. Reconcile these against the incident definitions used in Task 8.1 so that an internal incident and a reportable event are not assumed to be the same thing.
9.1.1 9.1.2 both TBD	Task 9.1 maps competency "from executive level to front-line operations," but both implementations are open and Tasks 9.2 and 9.3 are operator-facing. The executive entry in the competency matrix is the one most likely to be filled with general AI literacy content, which is not what the role requires.	Define the executive-level entry by the decisions the officer makes rather than by AI literacy in the abstract. See Task 13.7. General AI familiarity is not the competency at issue; the ability to test an assertion in a report is.

Index	Observation	Suggested change
11.1 11.1.1, 11.1.2 TBD	Ranking AI-dependent essential functions by impact on mission continuity and public safety allocates scarce recovery resources across competing public-safety priorities. That is an oversight decision presented as a technical management task.	Add: “The ranking should be ratified at the accountable level and reviewed on the cadence at which the organization reviews its continuity plan, since it determines the order in which public-safety functions are restored.”
Practices 1 to 12 collectively	No practice provides for assurance independent of the functions that operate and manage AI. Every control is first-line or second-line. The Design Criteria require outcomes that are audit-ready and that enable regulatory oversight, but nothing establishes a reviewer.	Adopt Task 13.6. Assurance placed under Practice 10 would sit inside the evidence-producing function it is meant to test, which is the defect rather than the remedy.
Appendix “Do no harm” design criterion	The criterion permits an operator to satisfy a control “to the letter” by formally declaring an operational path out of scope. The underlying concern about blind compliance and misinformed blanket mandates is well placed and the criterion should be kept. As written, however, a declaration requires no acceptor and no expiry, and is therefore indistinguishable from an undocumented gap. The criterion’s closing caution that “boundaries should be constructed tightly to prevent loopholes” addresses a different problem, and addresses it to the drafters: it governs what may be declared out of scope. It does not identify who accepts the consequence of a declaration, or when the declaration expires.	Add to the criterion: “Where the text permits an operator to declare an operational path out of scope, the profile should also require that the declaration be recorded with the name and role of the person accepting the residual risk, the compensating controls, and a review date.” This strengthens the criterion rather than narrowing it. It preserves the professional judgment the criterion exists to protect, and makes the exercise of that judgment attributable and time-bound.
5.1.1 observation	The Implementation requires cryptographically verifiable, non-shared identity for every autonomous sub-agent and transient tool invocation, on the ground that shared credentials defeat attribution. The reasoning is sound and, applied consistently, reaches the humans who authorize those agents to act on regulated processes.	No change to 5.1.1 is proposed. The observation is offered as support for Practice 13: the profile’s own attribution standard for machine actors is stricter than its attribution standard for human accountability.

Index	Observation	Suggested change
Appendix B incident mapping	The mapping is useful and the level of detail is right. It also shows the gap this comment describes, in the profile's own material. At i7, Uber ATG in Tempe, the incident maps to Practices 1, 3, 4, and 9, which between them reach the vehicle behavior, the handover, and the operator's automation complacency. The NTSB probable cause names those same failures and then attributes them: "Contributing to the crash were the Uber Advanced Technologies Group's (1) inadequate safety risk assessment procedures, (2) ineffective oversight of vehicle operators, and (3) lack of adequate mechanisms for addressing operators' automation complacency, all a consequence of its inadequate safety culture." Safety Recommendation H-19-52 to the company was to complete implementation of a safety management system comprising safety policy, safety risk management, safety assurance, and safety promotion. The mapped Practices capture the proximate failures. The organizational finding that the Board treated as their common cause has nowhere to go. At i3, Cruise, Practice 8 correctly reaches the incomplete disclosure as an incident-analysis defect. What no Practice reaches is who was accountable for deciding what to report.	Retain the mapping. Add a column recording what the official finding actually faulted, so a reader can see when the cited finding sits outside every mapped Practice, and map to Practice 13 where the finding was organizational. Without that, the appendix reads as though the twelve Practices exhaust the causes an official body identified, which for i7 they do not.

4 Responses to the questions put to reviewers

4.1 Glossary priorities

The question asks which terms are most likely to cause confusion when bridging AI development and critical infrastructure governance, and what semantic conflicts should be clarified. The most consequential conflicts are internal to the draft rather than between sectors.

Governance. The draft uses the word in at least three senses: enterprise risk governance in Section 1.1, the Practice layer described as governance-focused in Section 2, and governance and compliance personnel as an Implementation audience in the same section. A reader cannot tell from context which sense is meant. We recommend three entries: *oversight governance*, which determines what the organization is answerable for and who answers; *operational governance*, which determines how a system is run and by whom; and *technical governance*, which determines the controls embedded in the system itself. Defining each against the other two is more useful than defining any of them alone.

Oversight. Used in two unrelated senses. In Practices 3 and 4 and throughout Task 4.5 it means real-time supervision of a running system by a human operator. In the title of Practice 3 and in the Govern mappings it means accountability for the program. These are different activities, performed by different people, on time scales that differ by several orders of magnitude. We recommend either two entries or consistent qualification: *operational supervision* for the first sense, *oversight accountability* for the second.

Accountable. Table 2 defines the term as a system property. The governance sense, in which a named person is answerable for an outcome, does not appear. This is the single definition most likely to produce a

communication breakdown between an AI vendor and a regulated operator, because both will use the word sincerely and mean different things.

Authorization. Used at Implementation 12.1.3 for promoting a machine-authored artifact to production, and implied at Implementations 1.3.1 and 4.5.2 for the autonomy regime of the system itself. These are different decisions at different levels. The profile currently authorizes the artifact but not the system.

Risk tolerance, risk appetite, acceptable risk threshold, and safety margin. The draft uses risk tolerances at Sections 1.2 and 2 and in the Map 1.5 mappings at Practices 3 and 6, risk profile at Section 1.1, acceptable risk thresholds at Implementation 3.1.3, and safety margins at Implementation 8.2.2. In examined sectors these are distinct terms with distinct approval levels, and treating them as synonyms will cause a mapping error when an operator crosswalks the profile against existing sector requirements.

Owner. Section 1.1 uses owners in the critical infrastructure sense of asset ownership. The AI governance literature uses model owner to mean an accountable individual. Supply chain conversations that cross this boundary will collide, particularly where a vendor and an operator are negotiating the allocation of responsibility described in Section 1.2.

4.2 Document Design Criteria

The current criteria are sound and we would keep all of them. Three additions or amendments.

Add a locatable-owner criterion. For each Practice, Task, and Implementation, a reader should be able to determine which layer of the organization owns the decision the row describes, without reference to any particular organization's structure. Several rows presently describe a decision without indicating the level at which it is made, notably 3.1.3, 6.1.2, and 11.1. This criterion would catch that class of defect systematically rather than row by row.

Amend the “Do no harm” criterion as set out in Part 3 above, to require that an out-of-scope declaration name its acceptor and carry a review date.

On the “Achievable and provides a gradient” criterion, currently marked exploratory: keep it. The paired minimal and fuller implementations are the mechanism most likely to make this profile usable by a small water system or a rural cooperative, and abandoning them would concede the resource-constrained reader to another document. One qualification. The gradient should apply to the mechanism, not to the accountability. Whether an organization maintains its inventory in a governance platform or a spreadsheet is properly a matter of resources. Whether it has named an accountable officer is not: designation is among the least expensive controls in the profile and is precisely the control a resource-constrained operator most needs, because it determines who decides which of the twelve Practices that operator can afford to implement.

4.3 Prioritization from the Ideas and References Already in Progress list

The changelog records that review versus action velocity has been drafted as Task 4.5 and removed from this list. It would have been our first pick, so we start there. Our three from the current list share a property with it: each is framed as an engineering or human-factors consideration when it is in fact a decision that requires an owner.

On review versus action velocity, now Task 4.5. The task is well drafted and the core observation is correctly stated: human oversight degrades into a nominal check when the required review latency exceeds the system's physical action time. Implementation 4.5.1 documents the condition and Implementation 4.5.2 substitutes deterministic boundary constraints. What the task does not do is name who authorizes the substitution. Entering a human-out-of-the-loop regime is the point at which an organization ceases to have a human safety net and begins operating on deterministic boundaries alone. That transition should be

authorized, bounded, and re-ratified rather than inferred from a timing analysis. Implementation 13.4.2 supplies only that step and requires no change to 4.5 itself.

First, theoretical versus actual capability. Both the supply chain go/no-go decision at Implementation 6.1.2 and the autonomy determination at Implementation 1.3.1 turn on this distinction, and both are currently unowned. Making the distinction explicit without also naming who applies it would add precision to a judgment nobody is recorded as making.

Second, cost as a risk-management input. The note asks whether this deserves standalone treatment or should be folded into the existing gradient language. It deserves standalone treatment, and it should be located at the oversight layer. When a resource-constrained operator can implement three of twelve Practices, the selection of which three is a risk acceptance decision with public-safety consequences. Folding it into the gradient language recasts that decision as an engineering judgment about affordability, which is where it will be made by default and where it should not be made.

Third, autonomy is not one-dimensional. We agree, and we would offer the dimensions that matter for oversight purposes, which are not the same as those that matter for engineering. Four determine what must be authorized and at what level: the scope of authority, meaning what the system may act upon; the reversibility of the action; the time to consequence relative to human review time; and whether the action is observable to a human before it takes effect. A single autonomy level collapses all four, and the collapse is what produces the common failure in which a low-consequence system with broad, irreversible, unobservable authority is governed less carefully than a high-consequence system with narrow and reversible authority.

On “living off the land”: the item is correctly identified and its oversight-layer expression is worth noting. Where an AI system causes harm through intended, legitimate functionality rather than through malicious code, the operative control is not detection. It is the scope and expiry of the authority granted under Task 13.4. Authority granted broadly and permanently for a legitimate purpose is the vector.

5 Reference notes and additions

5.1 A scope caveat on the reference query at Implementation 2.3.4

The reviewer note at Implementation 2.3.4 states that the team is evaluating the relevance of OCC Bulletin 2026-13 and ISO/IEC 42001:2023 for determining evaluation cadence. The bulletin is current and correctly identified. The issue is its scope, and it is the kind of thing that is expensive to discover late.

OCC Bulletin 2026-13, Federal Reserve SR 26-2, and FDIC FIL-15-2026, all issued 17 April 2026, constitute the revised interagency guidance on model risk management. Bulletin 2026-13 rescinds OCC Bulletin 2011-12 among other issuances, and SR 26-2 supersedes SR 11-7 and SR 21-8. They are sound references for ongoing monitoring and periodic review cadence for models generally. A footnote to Section II, Purpose and Scope, at page 3 of the guidance, however, expressly excludes the systems this profile principally addresses: “Generative AI and agentic AI models are novel and rapidly evolving. As such, they are not within the scope of this guidance.” Citing the bulletin for AI evaluation cadence without that caveat will mislead financial-sector readers, who will apply a scope carve-out the profile did not intend to invoke. We recommend citing it with the caveat and noting that the agencies have signaled separate AI-specific guidance to follow.

ISO/IEC 42001:2023 is a reasonable cadence reference, but cadence is not the standard's principal utility to this profile. Clause 5, Leadership, and in particular 5.3, Roles, responsibilities and authorities, is the part of the standard that addresses the layer this profile does not currently reach. We would cite it there.

5.2 ISO/IEC 38507:2022 is absent from the draft entirely

ISO/IEC 38507:2022, *Governance implications of the use of artificial intelligence by organizations*, does not appear anywhere in the draft, including in the Standards Bodies, Consortia, and Efforts We're Tracking list in the Appendix. It is the one international standard written specifically for members of the governing body on the organizational use of AI, and Clause 4, Governance implications of the organizational use of AI, addresses the profile's gap directly. Clause 4.1 provides that the governing body remains accountable for all activities of an organization and that this accountability cannot be delegated. Clause 4.3, Maintaining accountability when introducing AI, provides that the governing body should take responsibility for the use of AI rather than attributing responsibility to the AI system itself, and that its accountability should be established across all aspects of intended or actual use of AI. That is the oversight layer stated in an international standard, and it is the most transferable reference available for a cross-sector profile.

We recommend adding it at Practice 13 if that Practice is adopted, and at minimum adding it to the tracking list alongside ISO/IEC TR 5469. We note that ISO/IEC 42001:2023 currently appears only in the reviewer note at Implementation 2.3.4 and is not itself on the tracking list.

5.3 Additions to the accountability references at Practice 13 and at 8.5

For the designation mechanism proposed at Implementation 13.2.1, the field-tested references are those already operating in the sectors themselves. NERC Reliability Standard CIP-003-9, Requirements R3 and R4, is the most fully developed: R3 requires identification of a CIP Senior Manager by name with any change documented within thirty calendar days, and R4 requires that delegations be documented by name or title, approved by the CIP Senior Manager, and updated within thirty days of change. The structure is instructive because top-level accountability is non-delegable while downstream authority is delegable under a documented and current record. The same language is carried forward unchanged in CIP-003-11, approved by FERC in Order No. 918 on 19 March 2026 and enforceable in 2029. The parallel mechanisms are 45 C.F.R. 164.308(a)(2) in healthcare, which requires identification of the security official responsible for the development and implementation of the required policies and procedures, and TSA Security Directive Pipeline-2021-01G, Sections II.B.1.a and II.B.2.b, in pipelines, which requires a primary and at least one alternate Cybersecurity Coordinator at the corporate level, accessible to TSA and CISA at all hours. That directive cancels and supersedes SD Pipeline-2021-01F and is effective 16 January 2026 through 15 January 2027. TSA's proposed *Enhancing Surface Cyber Risk Management* rule, 89 Fed. Reg. 88488 (7 November 2024), would additionally require identification of an "accountable executive" alongside the Cybersecurity Coordinator; the term is a useful one for this profile whether or not that rule is finalized.

For the assurance structure proposed at Task 13.6, the reference is *Assurance and Advice in Support of Effective Governance: Three Lines Model, Statement of Position* (Institute of Internal Auditors, 2026), which supersedes the July 2020 Three Lines Model and, before it, the 2013 Three Lines of Defense. Separately, the Institute has issued no Topical Requirement on AI. Its Topical Requirements issued to date address cybersecurity, third party, organizational behavior, and organizational resilience. Its Artificial Intelligence Auditing Framework is guidance rather than a mandatory element. Task 8.5 already states that the market for third-party AI forensic auditors is nascent and offers three interim approaches. The absence of an AI Topical Requirement is a concrete data point supporting that statement and worth citing in it, since it tells a reader how far the assurance profession has and has not moved.

5.4 Two small corrections

Section 2 refers to "Appendix A: Historic Incident-to-Practice Mapping." The incident mapping is Appendix B; Appendix A is the trustworthiness characteristics placeholder. Section 1.1 has it right.

The reviewer note preceding the table describes the Cyber AI Profile as "still-drafting," and Section 1.2 deconflicts with it without stating its status. For readers planning against it, the precise status as of this

submission is that the initial preliminary draft was released on 16 December 2025 with a comment period that closed on 30 January 2026, and an initial public draft has not yet been issued. Since the Primary Cyber AI Profile Mappings field appears throughout the table, readers will reasonably want to know how long those fields are likely to remain open.

5.5 Sector reference lists

The Electric and Bulk Power list in the Appendix cites NERC CIP-002 through CIP-014 with emphasis on CIP-013-2 for supply chain controls. We suggest calling out CIP-003 Requirements R3 and R4 specifically, since the accountability structure they establish is the element most transferable to the other fifteen sectors, whereas the supply chain content of CIP-013 is largely sector-specific.

6 About the Center for AI Oversight

The Center for AI Oversight is a 501(c)(3) nonprofit defining oversight governance for AI in regulated industries. Its work 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. Its Executive Director is lead author of the AI oversight body of knowledge published by the Private Directors Association, the national association of directors of privately held and family-held companies, and has comparable material underway with the DCRO Institute, which credentials directors and executives in risk governance.

Brian J. Allen, Esq., is Executive Director. He assessed and built cybersecurity oversight practices for organizations on a global scale at EY, represented the banking sector, through the Bank Policy Institute, in the development of the NIST AI Risk Management Framework, and was appointed by the Chairman of the Federal Communications Commission to represent the communications sector during the development of the NIST Cybersecurity Framework. The Center's advisory board is Jim Cunha (Federal Reserve Bank of Boston, ret.), Chuck Senatore (SEC Southeast Region, former), David Simpson (Rear Admiral USN ret.; Chief, FCC Public Safety and Homeland Security Bureau, former), Brian Peretti (CTO and Deputy Chief AI Officer, U.S. Treasury, ret.), and Ankur Singhal (senior risk and technology executive). Advisory board members provide strategic counsel and are not individually endorsing these recommendations.

The Center is available for a follow-up conversation on any part of this comment, and will draft further Practice 13 material, including implementations for rows currently left open, if that is useful to the team's backlog. Brian J. Allen, Esq., cfaio.org.

Disclosure of AI assistance. This comment was prepared with the assistance of an AI assistant (Anthropic Claude). The assistant was used to structure the argument, to draft and revise text against the author's positions, and to check the currency of regulatory and standards citations. The structural finding, the three-layer framing, the proposed Practice 13 content, and the judgments about which items to prioritize are the author's, drawn from direct work with operators, executives, and boards in regulated industries. All citations were verified against primary sources; where verification was incomplete, the claim was removed rather than qualified.

Status of citations. Regulatory and standards citations are stated as of 11 August 2026. Index numbers and quotations from the profile are keyed to the Community of Interest Discussion Draft of 3 August 2026. TSA's *Enhancing Surface Cyber Risk Management* rule (RIN 1652-AA74) is cited as proposed: comments closed 5 February 2025 and no final rule has issued.

A Appendix: the designation mechanism in each of the sixteen sectors

This table answers the question a cross-sector profile has to answer before routing to sector mechanisms: does the mechanism actually recur. It does in three quarters of the sectors and fails in a quarter. Instruments are current as of 11 August 2026.

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 accountable for high-impact AI use, independent review, and risk acceptance. The federal government has imposed on itself the named AI accountability requirement it has imposed on none of the sectors it regulates. That is the gap this Practice addresses, and it is why the cross-sector profile is the right instrument: the mechanism exists in most sectors and reaches AI in almost none of them. Routing to it is cheap. Waiting for sixteen regulators is not.

Read it as a map of available plumbing, not of where the duty lives. The duty arises from the governing body's obligation to oversee mission-critical risk and from the standard of care expected of a critical infrastructure operator. In the three sectors with no federal mechanism, the charter is issued on the governing body's own authority.

Sector	Instrument and provision	What it designates	Reaches AI?
Chemical	CFATS, 6 C.F.R. 27.230(a)(17) and 27.200(b)(3). Statutory authority lapsed July 2023. Enforceable backstop: EPA Risk Management Program, 40 C.F.R. 68.15(b)	Individual, but the in-force backstop permits designating a position rather than a person	Silent
Commercial Facilities	No federal designation requirement. CISA sector guidance is voluntary	Nothing	Silent
Communications	47 C.F.R. 1.20003(a) and (b) (CALEA); 47 C.F.R. 64.2009(e) (CPNI)	Individual, by name, with an annual officer certification filed with the FCC	Silent
Critical Manufacturing	No federal designation requirement. NIST IR 8183r1 and CISA guidance are voluntary	Nothing sector-wide	Silent
Dams	FERC, 18 C.F.R. 12.62(a), (b), (d)	Individual: a Chief Dam Safety Engineer or Coordinator, who signs the program	Silent
Defense Industrial Base	CMMC, 32 C.F.R. 170.22 and 170.4(b)	Individual: an Affirming Official, by name, filed in a government system	Silent
Emergency Services	OSHA HAZWOPER, 29 C.F.R. 1910.120(q)(3)	Individual, but per incident by operation of law, not an appointed officer	Silent

Sector	Instrument and provision	What it designates	Reaches AI?
Energy	NERC CIP-003-9, R3 and R4	Individual, by name, with a thirty-day change duty and documented delegations	Silent
Financial Services	GLBA Safeguards Rule, 16 C.F.R. 314.4(a) and (i); 12 C.F.R. Pt. 30 App. D, I.E.3 and II.C	Individual: a Qualified Individual reporting in writing to the board at least annually	Silent, technology-neutral
Food and Agriculture	FDA, 21 C.F.R. 117.180 and 117.310; FSIS, 9 C.F.R. 417.2(d)	A qualified role that must be filled, plus a role-based plan signature	Silent
Government Services and Facilities	FISMA, 44 U.S.C. 3554(a)(3); OMB M-25-21, s.3(a)	Individual: a Senior Agency Information Security Officer, and a Chief AI Officer	Yes, expressly
Healthcare and Public Health	HIPAA Security Rule, 45 C.F.R. 164.308(a)(2)	Individual: the security official responsible for development and implementation	Silent, technology-neutral
Information Technology	No federal designation requirement. Reg. S-K Item 106(c)(2) requires disclosure of whether a designation exists, not that one be made	Nothing	Silent
Nuclear	10 C.F.R. Pt. 55 and 50.54(i), (k), (m); cyber rule at 10 C.F.R. 73.54	Individual licensure of operators, the strongest personal accountability here, but attached to operating acts; the cyber rule is entity-only	Silent. NRC gap assessment ML24290A059 says its rules do not explicitly address AI
Transportation Systems	49 C.F.R. 1570.201 (surface); 1542.3 and 1544.215 (aviation); TSA SD Pipeline-2021-01G, II.B.1.a and II.B.2.b	Individual, by name, filed with the regulator. The cleanest example in the table	Silent
Water and Wastewater	AWIA 2018 s.2013, 42 U.S.C. 300i-2	Entity-level certification only. No named individual	Silent

Eight sectors carry a mandatory federal named-individual designation at the program or officer level. Four carry a partial or role-based analog. One is entity-level only. Three have none. Confidence is lower on two rows: the CMMC affirmation obligation may be affected by reported changes to Phase 2 implementation, and the 12 C.F.R. Part 30 Appendix D threshold is the subject of a pending proposal (Docket ID OCC-2025-0207).