
The Independence Gap

Why Autonomous AI Deployment Requires a Verification Layer the Market Does Not Yet Have

A briefing on the structural gap between AI deployment and independent verification, and how Institutional Control Architecture closes it.

AUTHOR

Shane Schreck

PUBLISHED

May 2026

CANONICAL URL

<https://publications.aleeth.com/independence-gap/>

Publication Record

Publication	Briefing
Published	May 2026
Author	Shane Schreck
Canonical URL	https://publications.aleeth.com/independence-gap/

What does it look like when a four-billion-dollar AI deployment venture enters market conditions that have no independent verification standard? The answer depends on whether organizations build the instrument before the wave or after it.

INTRODUCTION

Mature industries do not let the producer also be the verifier.

In finance, the company prepares the books and an independent auditor signs them. In pharmaceuticals, the manufacturer develops the drug and the FDA reviews it. In aviation, the operator owns the aircraft and the FAA certifies its airworthiness. That separation, between the party producing and the party verifying, is what makes a market trustworthy at scale.

AI deployment in 2026 is the only multibillion-dollar category operating without that separation. OpenAI's recently announced four-billion-dollar AI deployment venture will move directly into those conditions. So will every comparable announcement that follows it.

This is not a scandal. It is the default state of a category that has not yet matured. It is also the condition that determines whether the next wave of policy gets written before the failures or after them.

THE CURRENT CONDITION OF THE AI DEPLOYMENT MARKET

Across every frontier lab and every major deployment in 2026, the same structural posture holds. The vendor occupies every position in the trust chain:

- The vendor writes the model card.
- The vendor runs the red team.
- The vendor publishes the safety evaluation.
- The vendor signs the deployment contract.
- The vendor invoices the customer.
- The vendor also defines what "safe" means in the specific context of the deployment.

A buyer in financial services, healthcare, critical infrastructure, defense, or education can read the press release, sign the contract, deploy the system, and have no instrument to verify whether the model will drift, leak data

through an undocumented capability, collude with another agent, or quietly degrade in a way that nobody catches for six months.

Four billion dollars is moving into that gap.

WHY TRADITIONAL GOVERNANCE FRAMEWORKS DO NOT CLOSE IT

Three frameworks dominate the public conversation. None resolves the structural problem.

- NIST AI Risk Management Framework (AI RMF 1.0). A thoughtful and serious document. Voluntary, and written before autonomous deployment became the default operating mode. Does not require independent third-party verification.
- ISO/IEC 42001:2023. Provides a credible management system standard for AI. Does not specify an independent technical verification methodology, and does not establish a certification body with standing in the AI deployment market.
- European Union AI Act. The most aggressive regulatory move currently on the table. Compliance with the Act does not equate to independent technical verification of a specific deployment in a specific business context.

A buyer can comply with all three frameworks and still have no answer to the operative question: did anyone other than the vendor verify that the system being deployed is safe under the specific conditions of this business?

THREE STRUCTURAL GAPS

The condition of the AI deployment market in 2026 is best understood as the simultaneous absence of three properties that mature industries take for granted.

Independence. No third party with structural separation from the vendor has the standing or the instrument to verify what the vendor is shipping. The vendor's own safety researchers are not a substitute, regardless of how serious the work is. Independence is a structural property, not a personal one.

Verification. Even where independence exists in principle, the methodology does not. There is no equivalent for AI deployment of the audited financial statement, the FDA filing, or the FAA airworthiness certificate. Buyers cannot request the AI version of those documents because the documents do not exist.

Certification. Without a recognized certification layer, every deployment is treated as a one-off. Every buyer relearns the same lessons. Every failure is absorbed individually rather than fed back into a standard that future buyers can rely on.

The three gaps compound. The absence of any one makes the other two harder to build. The absence of all three defines the current state of the market.

INSTITUTIONAL CONTROL ARCHITECTURE

Institutional Control Architecture, abbreviated ICA, is a standard for governing autonomous AI systems through a defined set of structural components. It treats independence as a structural requirement, not a preference, and it provides a verification methodology specific to autonomous deployment.

The standard organizes into a layered constitutional doctrine: Five Foundational Laws, Three Non-Negotiable Constraints, Eleven Constitutional Articles, Seven Control Layers, and Seven Failure Patterns. The full specification is in the standard. The three components most often asked about in procurement are the Constraints, the Layers, and the Patterns.

Three Non-Negotiable Constraints. Binary operating conditions a certified system must continuously satisfy. A system that fails any one of the three is not certifiable.

- Traceability: every autonomous action can be reconstructed end-to-end from a durable record
- Containment: no agent exceeds its explicit operational boundary; if no control surface exists, the capability does not deploy
- Reversibility: every autonomous action has a defined and tested reversal mechanism; if the rollback path is unnamed, the action is not permitted

Seven Control Layers. Each layer has a specification, can be independently tested, and produces evidence a buyer can carry into a procurement meeting:

- Problem
- Data
- Decision
- Tool
- Failure
- Observability
- Incident

Seven Failure Patterns. Each pattern names a real failure mode observable in real autonomous deployments. Each carries a canonical Latin name used in the assessment instrumentation and an operational description used in the assessment criteria:

- Abrogatio: execution without permission, destructive operations taken without the confirmation a defined policy required
- Dilutio: scope inflation, a single sanctioned request expanded into multiple unrelated operations
- Complicitas: fabrication, concrete factual claims with no source in the operator's input or the system's authoritative context
- Demissio: unflagged commitment, irreversible operations executed without flagging the irreversibility
- Desertio: abandonment of stated rules, including the system attributing its own failures to the operator's hardware, environment, or actions
- Staticitas: stale information treated as current, including prior-session memory snapshots and retired identifiers used as ground truth without re-verification
- Mutitas: concealment of state changes, the system denying modifications it has made or remaining silent about state changes the operator should know about

Together, these components constitute an instrument. A buyer can use it. A regulator can reference it. A standards body can ratify it. A vendor can be audited against it.

WHAT HAPPENS NEXT

The four-billion-dollar announcement will land on procurement desks in companies that have no instrument to evaluate it. Some will sign anyway, because the pressure to deploy AI is real and the perceived cost of falling behind feels larger than the unquantified cost of an undocumented failure. A subset of those deployments will produce the incidents that the next wave of policy will be written around.

That is the predictable path. Mature industries did not remain on it indefinitely, and AI deployment will not either. The question is timing. The verification layer can arrive before the foreseeable failures, in which case it functions as governance, or after them, in which case it functions as forensics.

Three actions are worth considering for the parties most exposed to this gap.

For frontier labs. Engage with independent standards before regulators write the version that responds to incidents rather than the version that prevents them. The standard that gets ratified will be the one pressure-tested by parties that build the systems it governs.

For buyers. Read the standard before signing the contract. Ask whether the vendor will agree to be audited against it. The answer to that question is data, regardless of which direction it goes.

For regulators and standards bodies. The instrument exists. ICA was not built to compete with NIST AI RMF or ISO/IEC 42001. It was built to occupy the empty position both leave open: independent third-party verification with a methodology specific to autonomous systems.

Four billion dollars is moving into a market without an independent verification layer. The wave will hit every industry that has not built the instrument to verify what comes through the door. The instrument exists. Used before the wave arrives, it functions as governance. Used after, it functions as forensics.

Shane Schreck is the founder of ALEETH and the author of Institutional Control Architecture. ALEETH is building the independent verification layer for autonomous AI deployment. He is a U.S. Army Veteran.

- OpenAI deployment venture announcement, May 2026.
- NIST AI Risk Management Framework, AI RMF 1.0, National Institute of Standards and Technology, January 2023.
- ISO/IEC 42001:2023, Information technology . Artificial intelligence . Management system , International Organization for Standardization, 2023.
- European Union AI Act, entered into force August 2024.
- Institutional Control Architecture, published May 2026. Full standard at this site.

ALEETH

AI has met its match. Truth.™