25RevolAI

WHITE PAPER

The AI Never Told Me It Was Required

A postmortem of the first FDA enforcement action citing artificial intelligence, what the federal accountability standard now means for the private sector, and why the answer is not a black box.

July 15, 2026

Dave Stephens | Operating Partner, 25RevolAI
dave.stephens@25revol.com | +1 801 870 1449

© 2026 25RevolAI | 25revol.ai White Paper: The AI Never Told Me It Was Required
All rights reserved. Unauthorized reproduction is prohibited.

Executive Summary

On April 2, 2026, the FDA issued Warning Letter 320-26-58 to Purolea Cosmetics Lab, a small drug manufacturer in Livonia, Michigan. Most of the letter reads like a hundred others: insanitary conditions, missing release testing, a quality unit that was not doing its job. One section does not. For the first time, the agency cited a firm for inappropriate use of artificial intelligence in pharmaceutical manufacturing [4]. The firm had used AI agents to write its drug product specifications, its procedures, and its master production records. When investigators asked why process validation had never been performed before product went out the door, the owner gave an answer that should be read carefully by every executive currently evaluating AI: the AI agent never told her it was required.

The FDA did not need a new regulation to act. It applied 21 CFR 211.22(c), a quality unit requirement that has been on the books for decades. That is the first lesson of this paper: the accountability standard for AI in regulated operations is not coming. It is already here, and it is being enforced with existing rules.

The federal government has been converging on this standard for years. In 2021 the Government Accountability Office published its AI Accountability Framework, GAO-21-519SP, built on four principles: governance, data, performance, and monitoring. GAO addressed it explicitly to federal agencies and other entities, meaning you. In 04/2026, GAO-26-107859 found that federal agencies more than doubled their use of AI in a single year while failing to systematically capture the lessons learned, best practices, and testing requirements that responsible AI acquisition depends on. Regulators are not waiting for perfect guidance, and they are not giving the private sector that luxury either.

The second lesson matters more. Purolea did not fail because it used AI. It failed because it used a black box: a general-purpose model that generates plausible text, did not say where that text came from, cannot describe what it does not know, and stays silent about both. Every gap the tool had, the firm inherited without knowing it. This paper walks through what happened, why it happened, how the four federal accountability pillars translate to private sector operations, and what an architecture looks like when accountability is designed in rather than hoped for.

1. What Happened: A Postmortem of Warning Letter 320-26-58

FDA investigators inspected the Purolea facility from Oct. 28-30, 2025. The firm manufactured homeopathic drug products, including topical products labeled for shingles and genital herpes relief, along with cosmetics. The inspection produced a Form 483, the firm's response was found inadequate, and the warning letter followed on April 2, 2026 [1].

All rights reserved. Unauthorized reproduction is prohibited.

The conventional findings

Set the AI section aside for a moment and the letter describes a facility with fundamental control failures:

- **Insanitary conditions.** Insects, filth, leaves, and clutter in manufacturing areas, and a docking bay door that exposed manufacturing directly to the outside environment when opened.
- **No release testing.** Finished drug products released without microbiological testing, a violation of 21 CFR 211.165(b). The firm had no scientific evidence its batches were free of objectionable contamination.
- **Untested components.** Components used without identity testing, relying on supplier certificates of analysis that were never qualified, violating 21 CFR 211.84.
- **Quality unit failure.** A quality unit that failed to establish procedures, review batch records before release, or establish production and process controls, violating 21 CFR 211.22.
- **Unapproved new drugs.** Products making disease treatment claims without approved applications, violating sections 301(d) and 505(a) of the FD&C Act.

The finding that makes this letter historic

During the inspection, the firm told investigators it had used AI agents to help it comply with FDA regulations. Specifically, the AI created the drug product specifications, the procedures, and the master production or control records. These are not peripheral documents. They are the spine of CGMP compliance, the documents every batch is manufactured against and every investigator reads first.

> *If you use AI as an aid in document creation, you must review the AI generated documents to ensure they were accurate and actually compliant with CGMP. Your failure to do so is a violation of 21 CFR 211.22(c).*
>
> FDA Warning Letter 320-26-58, 04/02/2026

Then came the admission that gives this paper its title. Investigators found the firm had never conducted process validation before distributing drug products, as 21 CFR 211.100 requires. The owner replied that she was not aware of the legal requirement, because the AI agent she used never told her it was required.

The FDA's prescribed remedy is worth quoting in substance: if the firm resumes production and uses AI for CGMP activities, any output or recommendation from an AI agent must be reviewed and cleared by an authorized human representative of the firm's quality unit. The agency did not ban the technology. It named the missing ingredient: an accountable human, backed by a system that makes review actually possible.

© 2026 25RevolAI | 25revol.ai White Paper: The AI Never Told Me It Was Required
All rights reserved. Unauthorized reproduction is prohibited.

2. Beyond the Surface: Why This Failure Is Systemic

The easy reading of this letter is that a small, sloppy shop cut corners. That reading is comfortable and wrong. Strip away the insects and the docking bay, and look at the decision the owner actually made: she asked a general-purpose AI agent to make the company compliant, received confident, professional-looking documents in return, and concluded the job was done. That is not an eccentric mistake. It is precisely what a large share of the market is being sold right now, wrapped in enterprise pricing. The failure has three root causes, and none of them is unique to a small business.

Root cause 1: Structural Fidelity v. Operational Truth

A general-purpose language model is trained to continue patterns. Ask it for a master production record and it will produce something that looks exactly like a master production record, because looking like one is what the model optimizes for. Whether the specification limits are right, whether the test methods exist, whether the procedure matches your actual process: the model on its own has no mechanism that checks any of it. The documents at Purolea did not look wrong. That was the problem.

Root cause 2: A black box cannot describe its own boundary

The most instructive detail in the letter is what the AI did not say. It never mentioned process validation. That silence was not a malfunction. A probabilistic model has no representation of the boundary of its own knowledge, so it cannot warn you when you have walked past it. It will answer the question you asked and stay silent about the ten questions you did not know to ask. At Purolea, one of those silent questions was a federal requirement, and the gap stayed invisible until an investigator walked in.

Root cause 3: The tool quietly became the quality unit

21 CFR 211.22 assigns responsibility to a quality unit made of accountable people. At Purolea, no named human owned the AI's outputs. Nobody reviewed the documents against the regulations, and nobody could have traced why the documents said what they said, because the model offers no provenance. The black box failed in both directions: no visibility into why an output exists, and no visibility into what is missing from it.

The enterprise translation

Purolea's blast radius was small. The same pattern at enterprise scale looks like this: an AI module bolted onto a legacy quality platform to make a brittle system feel modern, or a compliance copilot built by injecting regulatory text into a general-purpose model after the fact and hoping retrieval covers the gaps. We call the first one paid camouflage and the second one post hoc knowledge injection. They are the same sin at two layers of the stack: compliance as an afterthought instead

All rights reserved. Unauthorized reproduction is prohibited.

of a foundation. The difference between Purolea and an enterprise running these tools is not the failure mode. It is the number of batches released before anyone notices, and the size of the remediation bill when someone finally does.

3. The Federal Standard Arrived Without a New Rule

Executives waiting for an FDA AI regulation before taking this seriously have misread the situation twice.

First, the accountability framework already exists. GAO published GAO-21-519SP in 06/2021, organizing responsible AI use around four principles: governance, data, performance, and monitoring [2]. The framework is written for federal agencies and other entities, and it was built with input from industry precisely because GAO expected it to travel. Companies in the federal supply chain are measured against it directly. But the more important point for every regulated manufacturer is simpler: these four pillars are just what accountable AI looks like, and your own regulator does not need GAO's framework to hold you to them. The FDA just proved that.

Second, enforcement does not need new rules. GAO-26-107859, published 04/13/2026, examined how federal agencies acquire AI and found that even as they more than doubled their use of it, they were not systematically collecting the lessons learned or best practices, such as contract terms on data rights and testing, that sound acquisition requires [3]. The gaps GAO found in federal buying, no lessons captured on testing, on data rights, on what a vendor's system can actually do, are the same gaps a private quality leader inherits the moment she buys a compliance copilot on a vendor's word. Read that carefully: the government is publicly auditing itself for AI governance gaps while agencies keep adopting the technology. Nobody is waiting for perfect guidance. The FDA proved the same point at Purolea by enforcing AI accountability with a regulation written decades before the technology existed. The obligation was always there. AI just gave firms a new way to violate it.

Here is how the four federal pillars translate to a regulated private sector operation, and where Purolea failed each one:

Pillar	What accountability requires	Where Purolea failed	What you must be able to show
Governance	A named, authorized human owns every AI output that touches a regulated decision.	The AI became the de facto quality unit. Nobody reviewed its documents.	Documented quality unit review and clearance of every AI-generated output, per 21 CFR 211.22(c).
Data	You know what the system knows, and where every statement came from.	Generic model, unknown training data, zero provenance for any document.	Authoritative source grounding with traceable provenance for every statement the system makes.
Performance	Evidence the system does what is claimed,	No validation of the process or the	Deterministic, reproducible outputs with

© 2026 25RevolAI | 25revol.ai White Paper: The AI Never Told Me It Was Required
All rights reserved. Unauthorized reproduction is prohibited.

	benchmarked and documented.	documents that defined it.	validation and audit-ready documentation.
Monitoring	Failures and omissions are detectable and escalated during operation.	The validation gap was invisible until an FDA investigator found it.	Explicit knowledge boundaries that make omissions visible, plus a complete audit trail.

4. A Different Way: Accountability as Architecture

The market's dominant answer to the problems above is a disclaimer: use AI, but check its work. That advice is correct and insufficient. A quality unit cannot meaningfully review output it cannot trace, from a system whose knowledge boundary nobody can state. Telling operators to check a black box harder is not a control. The control must live in the architecture.

At 25RevolAI, we build for regulated industries on a different premise. Most vendors practice what we call Dependency as a Service™: bolt an AI module onto a brittle platform, inject compliance rules into a generic model after the fact, price it per seat, and rent the customer the risk. For a quality leader the honest name for that offering is Compliance Debt as a Service™. Our model is the opposite, Enablement as a Service™: domain expertise embedded in the foundation, an ownership model instead of a subscription trap, and scope-based billing instead of per-seat extraction. The architecture has three layers, and each one answers a Purolea failure directly.

Ground Truth Graph™: foundational knowledge alignment

The Ground Truth Graph™ is a proprietary knowledge base built from the ground up on authoritative regulatory sources: FDA regulations and guidance, ISO standards, GxP requirements. It was built before the AI, not bolted on after. Every statement the system makes traces to a source an auditor can open. At Purolea, nobody could say where the documents came from. Here, provenance is native. That is the data pillar, answered at the foundation.

Ground Truth OS™: boundaries enforced by the platform

Ground Truth OS™ is the platform that makes the model operational: sovereign, fully isolated client environments with regulatory knowledge boundaries enforced at the architecture level. It is platform agnostic and runs today on sovereign private cloud infrastructure. These boundaries are enforced as native system properties rather than passive policy documents.

Source Constrained AI™: the silence problem, solved

By design, the AI operates inside the boundaries of the Ground Truth Graph™, so output is grounded in authoritative sources rather than statistical hope. We call this Source Constrained AI™. The contrast with Purolea is exact. There, the model's silence hid a federal requirement, and no one could interrogate what the system did not know. Under Source Constrained AI™ the knowledge

All rights reserved. Unauthorized reproduction is prohibited.

boundary is explicit and inspectable: if the source does not support a statement, the system does not make it, and the question of what the system does not know has an answer an auditor can verify. That is the monitoring pillar built into the performance pillar, instead of a report written after the failure.

The Informed Operator: the human the FDA asked for

The FDA's remedy at Purolea was an authorized human of the quality unit reviewing and clearing every AI output. That is not an add-on to our model. It is the model. The system hands work back to informed operators with best practice built in and sources attached, so review is a real act rather than a rubber stamp. Humans stay accountable. The AI stays constrained. That is the governance pillar, and it is also just 21 CFR 211.22(c) taken seriously.

One caution we state plainly because the industry too often will not: nobody can sell you a validated tool, and you should walk away from anyone who claims to. What we deliver is a system built to be validation and audit ready, with deterministic behavior, documented performance, and the audit trails to prove it in an inspection. Validation is something your quality system does. Our job is to make it possible instead of painful.

5. What to Do Now

Whether or not you ever touch our platform, five actions follow directly from the Purolea letter and the federal framework. Any regulated manufacturer using or evaluating AI should be able to close this list within a quarter.

- **Inventory.** Find every place AI touches a GxP document or decision, including the tools employees use informally. Purolea's exposure started with a tool nobody controlled.
- **Ownership.** Assign a named, authorized member of the quality unit who reviews and clears each AI output, in writing, to maintain immediate compliance with 21 CFR 211.22(c).
- **Provenance.** For any AI-generated content in your quality system, demand to see the sources that support it. If your vendor cannot answer, that is your answer.
- **Boundaries.** Ask every AI vendor to state what their system does not know and how omissions surface. A system that cannot describe its boundary cannot be monitored, and silence is how Purolea got hurt.
- **Evidence.** Keep performance documentation benchmarked against the claims you relied on when you bought the system, and an audit trail from source to output. This is what GAO's performance and monitoring pillars look like on the ground.

The Closing Argument

The Purolea letter will be cited for years as the moment the FDA first named AI in an enforcement action. The wrong lesson is to slow down on AI. Regulated manufacturers who sit this out will pay for caution in cycle time, headcount, and cost, and the agency itself was careful to say the

© 2026 25RevolAI | 25revol.ai White Paper: The AI Never Told Me It Was Required

All rights reserved. Unauthorized reproduction is prohibited.

technology is not the violation. The right lesson is that accountability is not a document you write after the fact. It is a property of the architecture you choose before the first output is ever generated. Choose one where the answer to the question that sank Purolea (i.e. “What did the system not tell you?”) can actually be known.

We make that test easy to run. Bring us one real use case from your operation and we will deliver a working, human-reviewed prototype against it in 24 to 48 hours at no cost. We call the reason this offer exists the Zero Imagination Gap™: you should never have to buy an AI system on faith, and with the right architecture you do not have to.

Dave Stephens | Operating Partner, 25RevolAI

dave.stephens@25revol.com | +1 801 870 1449 | www.25revol.ai

A note on how this paper was made: it was drafted with AI grounded in the primary sources cited below, then reviewed, verified, and cleared by the author. That is the Informed Operator model, practiced.

Sources

[1] U.S. Food and Drug Administration, Warning Letter 320-26-58, Purolea Cosmetics Lab, MARCS-CMS 722591, 04/02/2026. https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/purolea-cosmetics-lab-722591-04022026

[2] U.S. Government Accountability Office, Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities, GAO-21-519SP, 06/2021. https://www.gao.gov/products/gao-21-519sp

[3] U.S. Government Accountability Office, Artificial Intelligence Acquisitions: Agencies Should Collect and Apply Lessons Learned to Improve Future Procurements, GAO-26-107859, 04/13/2026. https://www.gao.gov/products/gao-26-107859

[4] Regulatory Affairs Professionals Society, FDA warns firm for inappropriate use of AI in drug manufacturing, Regulatory Focus, 2026. https://www.raps.org/resource/fda-warns-firm-for-inappropriate-use-of-ai-in-drug-manufacturing.html

[5] 21 CFR Part 211, Current Good Manufacturing Practice for Finished Pharmaceuticals, including 211.22, 211.84, 211.100, and 211.165.

© 2026 25RevolAI | 25revol.ai White Paper: The AI Never Told Me It Was Required
All rights reserved. Unauthorized reproduction is prohibited.