THE FDA DOESN'T CARE THAT YOUR AI WAS CONFIDENT

For manufacturers in regulated markets: how to put AI to work without flunking your next audit

For operations and quality leaders at manufacturers subject to FDA, ISO, or customer quality audits

There's a version of AI adoption that ends with a plant manager explaining to an auditor that a batch record was changed by "the system," that nobody approved it, and that there's no log of why. Every quality professional reading this just felt their stomach drop — because they know the finding writes itself.

AI vendors are selling speed. Your auditor buys **evidence**. Here's how to get the first without losing the second.

THE RULE THAT DOESN'T CHANGE

Regulated manufacturing already answered the AI governance question decades ago — it just used different words. Whether it's 21 CFR Part 11, ISO 9001, or your biggest customer's supplier quality agreement, the spine is identical:

> **Every record that matters must show what changed, who changed it, when, under what authority — and that record must be trustworthy.**

Notice the rule never asks whether "who" is a person or a machine. It asks whether you can *prove* it. AI doesn't get an exemption because it's impressive. A confident wrong answer with no audit trail isn't innovation; it's a 483 observation with better marketing.

THE 5 REQUIREMENTS FOR AI IN A REGULATED PLANT

Put these in front of any vendor — or your own team — before AI touches production, quality, or inventory records:

1. Attributable actions. Every AI action is logged under its own identity — not a shared system account — with a timestamp and what changed. If the AI's work blends into a generic "system update," you cannot defend the record.

2. Bounded authority. The AI's permissions are enforced by the system, not by policy documents. It *cannot* alter a released batch record, a spec, or a training record — the same way a line operator can't approve their own deviation.

3. Human sign-off where a signature lives today. Anywhere your SOPs require a signature — deviations, releases, changes — the AI drafts and a qualified person approves. The signature stays human. The paperwork gets faster; the accountability stays put.

4. The log survives the vendor. Audit trails belong in *your* systems, exportable, retained per your schedule. An audit trail that lives inside a SaaS subscription is evidence you're renting.

5. A demonstrated off switch. When something looks wrong, you halt the AI immediately — including queued work — and can show an auditor when and how you did. "We turned it off pending investigation" is a sentence auditors respect. Make sure it's true before you need it.

WHERE AI ACTUALLY EARNS ITS KEEP (ONCE GOVERNED)

The unglamorous, high-volume, evidence-heavy work: drafting deviation write-ups from the data already in the system for QA review · watching supplier certs, calibrations, and training records for expirations before they bite · reconciling receiving paperwork against POs and flagging only the exceptions · assembling the audit-response binder that currently eats a week of someone's quarter.

The pattern: **AI does the assembly, people do the signatures, the log does the proving.** That combination doesn't just survive audits — it makes audit prep the easiest it's ever been, because for once, the evidence collected itself.

Slingr builds custom operational software for manufacturers where this is the architecture, not an add-on: AI agents with their own identities and scoped permissions, human approvals where signatures live, and audit trails in software you own outright. `slingr.io`