

APTRIA.NET

Any process, automated. · aptria.net

Ben

Former Controller

aptria.net

Book a free call: aptria.net/book

FREE PRACTITIONER CHECKLIST

Will Your AI Pass an Audit? A 12-Point Checklist for Finance Teams

Every AI demo I've ever seen looks great in the conference room. The invoices route, the approvals fire, the data flows. Then the first exception hits — or worse, an auditor walks in — and the whole thing unravels because nobody thought about who was accountable when the AI got it wrong.

This checklist isn't theory. It's the 12 questions I'd ask about any AI automation before I'd let it near a finance process. Run through it against your current setup. If you fail three or more, the system will work exactly as long as nothing goes wrong — which is not a system, it's a demo you haven't tested yet.

How to use this: Go through each item against your actual AI setup — not the vendor demo, your production environment. Mark each as PASS or FAIL. Three or more FAILs is a red flag that needs to be fixed before you put the system in front of real money or a real auditor.

The Controller's Checklist — 12 Points

Each item has a question, what PASSING looks like, and what FAILING looks like. Be honest with yourself.

1 Can the AI mark its own work as complete and trigger a payment or release a deliverable?

PASS

An independent check — separate code, a separate system, or a human — must verify the work before anything is released or paid. The AI is never the final authority on its own output.

FAIL

The AI processes the invoice, marks it approved, and the payment runs. The worker is also the grader. This is the single most common way AI automation fails an audit — and a scammer's favorite door.

2 Does anything risky stop automatically for a human to review before it continues?

PASS

There are defined dollar thresholds, exception categories, or risk signals that hard-stop the process and require a named human to review and approve before the workflow resumes. Those thresholds are in the code, not just the documentation.

FAIL

The AI handles everything — with a note somewhere in the settings that says "require approval above \$10,000" that nobody has tested. Or the threshold exists but a smart enough edge case routes around it.

3 Does every step write to a log that nobody — including admins — can quietly edit or delete?

PASS

Every action (AI decision, human approval, exception, override) is written to an append-only or cryptographically signed log. If someone modifies or deletes an entry, the tamper is detectable. An auditor could reconstruct the full sequence from the log alone.

FAIL

There's a database table with timestamps, but anyone with admin access can UPDATE or DELETE rows. Or the log is a text file with no integrity check. Or there's no log at all — just the end result.

4 Can you reconstruct, step by step, exactly what the AI did on any given transaction — six months later?

PASS

Given a transaction ID, you can pull up the full decision trail: what data the AI saw, what rules it applied, what it concluded, who (or what) reviewed it, and when each step happened. This exists as a report, not a mental reconstruction.

FAIL

You can see the final outcome but not the reasoning path. Or the detail logs are only kept for 30 days. Or the reconstruction requires asking the engineer who built it to remember how it worked.

5 Is there a clear, documented answer to "who approved this?" for every high-risk AI decision?

PASS

For any decision above your risk threshold, there is a named human approver in the log — not "the system," not "auto-approved," not a generic role. A specific person is on record. That person can be called and asked about it.

FAIL

High-risk decisions are logged as approved by "AI Workflow" or "Automated Process." When an auditor asks who authorized a \$47,000 payment, the answer is "the system did." That answer does not hold up.

6 Does the system verify the data the AI used — not just the AI's conclusion?

PASS

The source data (vendor record, bank account, invoice amount, GL code) is validated against a trusted reference at the time of processing — not assumed to be correct because it was in the system. If the source data changed, that change is flagged.

FAIL

The AI processes whatever is in the database and trusts it. If someone updated the vendor's bank account number three minutes ago, the AI pays the new account without noticing. This is the exact vector for bank-change invoice fraud.

7 Does the system catch and handle exceptions — or does it just silently stop or pass them through?

PASS

When the AI encounters something it wasn't trained on or that falls outside its rules, it stops, flags the exception explicitly, routes it to a human, and waits. The exception is logged. Nothing proceeds until it is resolved.

FAIL

Unusual inputs are either ignored (the AI skips them and moves on), or worse, the AI makes its best guess and keeps going. Exceptions are the moments where AI is most likely to be wrong — and most likely to do real damage.

8 Are all human overrides of the AI's decision logged, with a reason and a name?

PASS

Every time a human overrides what the AI decided — approves something it flagged, rejects something it approved, or manually changes an AI-set value — that action is logged with the person's name, the timestamp, and a required reason field.

FAIL

Humans can override AI decisions with a button click and no record of who did it or why. In an audit, the log shows the AI flagged something, and then the outcome changed — with no explanation. That gap is worse than no flag at all.

9 Are the AI's duties separated from the authorization duties — the same way you'd separate duties for a human employee?

PASS

The AI can initiate or process, but it cannot also authorize. An independent role — human or a separate verification system — handles authorization. No single AI agent touches an entire transaction from input to payment without a checkpoint.

FAIL

The AI reads the invoice, codes it, approves it, and queues the payment. That is one entity controlling the full financial cycle. You would never let a single human employee do all four of those things without a review. The same rule applies to AI.

10 Can an independent auditor understand what the AI did — without you having to explain it?

PASS

The AI's decisions are logged in plain, human-readable terms: "Invoice flagged: bank account on record does not match vendor master. Held for human review. Reviewer: J. Smith. Approved: Yes. Reason: confirmed directly with vendor by phone." An auditor could read this with no technical help.

FAIL

The log contains model scores, vector distances, or confidence thresholds that require an engineer to interpret. Or the only record is a database row with a status code. Explainability is not optional — it is what "auditable" means.

11 Is the AI's "done" signal actually independent from its own processing — not just a self-report?

PASS

A separate verification step — outside the AI's own control — confirms completion before the downstream action fires. This could be a hash check, a second system querying a known-good state, or a human sign-off. The AI cannot unilaterally declare itself done and trigger consequences.

FAIL

The AI sets a status field to "COMPLETE" and the payment job polls that field. The AI is both the worker and the source of truth for its own work quality. This means a bug, a hallucination, or a compromised AI can release money by simply setting a flag.

12 Has anyone actually tried to break the system — run a fake invoice, change a bank account, corrupt the log — and documented what happened?

PASS

There is a documented test where someone intentionally ran a fraudulent input, a tampered record, or an out-of-bounds case against the live system. The result — what the system caught, what it missed, how it responded — is on record. It has been tested adversarially, not just optimistically.

FAIL

The only tests run against the system were success cases — real or realistic invoices that were supposed to go through. Nobody has run a fake bank-change invoice through it. Nobody has tried to edit a log entry. You do not know if the controls work because you have never tested them under adversarial conditions.

Scoring: If your AI automation fails 3 or more of these items, it will work in the demo and fail in the audit. That is not a software problem — it is a system design problem, and it needs to be fixed before the system touches real money or a real auditor.

The Three Controls That Cannot Be Missing

Every item on this checklist matters, but items 1, 2, and 3 are load-bearing. If those three are wrong, the rest doesn't matter:

- The AI's "done" must not trigger a payment by itself. An independent check has to pass first.
- Anything risky stops for a human — and that stop is in the code, not just the policy document.
- Every step writes to a log nobody can quietly edit.

Those three are not optional. They are the minimum. Everything else on this list builds on them.

Failed 3 or more? Let's talk about what it takes to fix it.

I do a free intro call — no pitch, no pressure. You walk away knowing exactly which of the 12 items your system is missing and what fixing them actually requires. I've built these controls into production systems; I know what it takes.

Book a free intro call: aptria.net/book · Or email directly: raining6693@gmail.com · 385-225-1199

Practitioner note: This is practitioner judgment as a former Controller — not legal advice or a certified framework. Take it to your auditor for a second opinion. The checklist reflects how I evaluate AI automation for finance teams, drawn from 8 years of Controller experience and direct work building verification systems. It is not an AICPA standard, a SOX requirement, or a regulatory opinion.