

AI Code Provenance Audit

Everyone used AI. That's not the finding. The finding is whether anyone understood what came out, and whether the person who has to fix it at 2am is the same person who accepted it at 2pm.

I call the gap comprehension debt. It's not bad code. Bad code you can read and repair. This is working code that nobody on the team can explain, which behaves correctly until the day it doesn't.

Run this on its own, or as the section most due diligence checklists are still missing.

The one question. Pick a non-trivial module the team says was AI-assisted. Ask an engineer to walk you through it line by line, out loud, without opening the assistant. What happens next tells you more than any metric on this page.

Scope: how much, and where

OK CON N/A

- ☐ ☐ ☐ Roughly what share of the current codebase was AI-generated or AI-assisted?
- ☐ ☐ ☐ Is the answer measured, or a guess?
 - A guess is fine. A confident guess with no way to check isn't.
- ☐ ☐ ☐ Which parts: scaffolding, business logic, tests, infrastructure, migrations?
- ☐ ☐ ☐ Is generated code concentrated in one area or spread everywhere?
- ☐ ☐ ☐ Was any of it generated by someone no longer on the team?
- ☐ ☐ ☐ Is there generated code in the parts that move money, touch customer data, or control access?

Tools and workflow

OK CON N/A

- ☐ ☐ ☐ Which assistants and agents are in use, and which were used historically?
- ☐ ☐ ☐ Are those tools still in the workflow, or was there a burst that nobody maintains?
- ☐ ☐ ☐ Does anyone pin model or tool versions for reproducibility?
- ☐ ☐ ☐ Are prompts, system prompts and agent configs in version control?
- ☐ ☐ ☐ Is there a written standard for when AI use is appropriate and when it isn't?
- ☐ ☐ ☐ Do engineers disclose AI assistance in pull requests?

Review and accountability

OK CON N/A

- ☐ ☐ ☐ Who reviewed generated code before it merged?

- ☐ ☐ ☐ Is that review recorded, or asserted?
- ☐ ☐ ☐ Did review get lighter as volume went up?
—This is the common failure. Throughput rises, attention doesn't.
- ☐ ☐ ☐ Are there commits that went in without a human reader?
- ☐ ☐ ☐ Can they show a pull request where a reviewer rejected generated code and why?

Comprehension

OK CON N/A

- ☐ ☐ ☐ Can a named engineer explain each AI-heavy module without the assistant open?
- ☐ ☐ ☐ If that engineer left tomorrow, who is next?
- ☐ ☐ ☐ Has anyone debugged a production incident inside generated code?
—The best possible answer here is yes, with a story
- ☐ ☐ ☐ Are there areas the team avoids touching because they aren't sure how they work?
- ☐ ☐ ☐ When something breaks in generated code, is the instinct to read it or to regenerate it?
- ☐ ☐ ☐ Could the team rebuild a given module from the requirements if the code vanished?

Tests and verification

OK CON N/A

- ☐ ☐ ☐ Are there tests around the generated code, and did a human write them?
- ☐ ☐ ☐ Do those tests assert behaviour, or do they restate the implementation?
- ☐ ☐ ☐ Has anyone deliberately broken generated code to confirm the tests fail?
- ☐ ☐ ☐ Are generated tests counted in the coverage number quoted to investors?
- ☐ ☐ ☐ Is there any property, fuzz or integration testing on generated logic?

Licensing and provenance

OK CON N/A

- ☐ ☐ ☐ Any generated code reproducing training material under an incompatible licence?
- ☐ ☐ ☐ Does the assistant's terms of service claim anything about output ownership?
- ☐ ☐ ☐ Is generated code marked in any way, in commits or comments?
- ☐ ☐ ☐ Has a licence scan been run since AI-assisted development started?
—If the last scan predates the tooling, the scan isn't evidence

Models in the product

Skip this section if no model runs at request time.

OK CON N/A

- ☐ ☐ ☐ Which models does the product call, from which providers?
- ☐ ☐ ☐ What happens on provider outage: degrade, queue, or fail?
- ☐ ☐ ☐ What happens when the model is deprecated? Providers retire models on their schedule, not yours.
- ☐ ☐ ☐ Is there a tested fallback to another provider or a smaller local model?
- ☐ ☐ ☐ Are model version, temperature and parameters pinned and recorded?
- ☐ ☐ ☐ What is the cost per request, and how does it move with usage?
- ☐ ☐ ☐ Is there a ceiling on spend, or can a traffic spike become a bill?
- ☐ ☐ ☐ Is there validation on model output before it reaches a user or a database?
- ☐ ☐ ☐ What is the plan for prompt injection, if user input reaches a prompt?
- ☐ ☐ ☐ Is there an evaluation suite, and does it run on every model change?

Training and fine-tuning

Skip this section if they don't train.

OK CON N/A

- ☐ ☐ ☐ Where did the training data come from, and can they prove the right to use it?
—Training data rights now appear as representations in purchase agreements. Provenance nobody can trace is a liability that scales with the dataset.
- ☐ ☐ ☐ Is data licensed, consented, scraped, or bought, and is that recorded per source?
- ☐ ☐ ☐ Do model cards and a training data datasheet exist, with known limitations?
- ☐ ☐ ☐ Are model artifacts versioned and reproducible from a recorded pipeline?
- ☐ ☐ ☐ Is there any signing or attestation on model artifacts?
- ☐ ☐ ☐ Any regulatory exposure identified, EU AI Act or sector rules?
- ☐ ☐ ☐ If a data source demanded removal tomorrow, what would that cost?

The call

OK CON N/A

- ☐ ☐ ☐ Is the AI use here leverage, or is it a loan against future maintenance?
- ☐ ☐ ☐ What is the cost, in weeks, of a team taking real ownership of the generated parts?
- ☐ ☐ ☐ Would you fund this team to maintain it, or only to have shipped it?