

PREPARED FOR

Sample Engagement

ILLUSTRATIVE EXAMPLE

PREPARED BY

Collin Wilkins

COLLINWILKINS.COM

ENGINEERING DIAGNOSTICS · EDITION I

# AI Readiness *Assessment.*

*A five-day diagnostic of how your engineering team adopts, owns, and measures the substrate underneath its AI tooling.*

## WHAT YOU GET

Written assessment specific to your team's stack, repo, and workflow — surfaced from your call and a read of your codebase.

12-14 PP.  
PDF

Custom AI Adoption Playbook adapted to your stack — sequencing, ownership, and measurement, scoped to your team size.

PLAIN  
ENGLISH  
PDF

Ready-to-commit **AGENTS.md** for your main repo — evaluable by your VP Eng in 90 seconds.

MARKDOWN  
REPO-READY

# YOUR PROBLEM AREAS

Engineering teams adopt AI coding tools faster than the substrate to use them well. Output looks productive at the developer level while velocity stays flat at the team level. Four specific gaps create that friction. Each is fixable, and none gets fixed by buying another seat.

01

## Convention files exist but aren't load-bearing

A `CLAUDE.md` lives in your main repo. It was created six months ago, last updated 11 weeks ago, and most of the team has forgotten it's there. Other repos have nothing.

Every Claude Code or Cursor session starts from a different context, output quality varies team to team, and onboarding a new engineer to "how we use AI here" depends on which senior they sit next to.

02

## AI tooling is owned by everyone and no one

One engineer is informally "the AI person." It isn't on her job description. She answers questions in Slack on her own time, and when she takes PTO the questions stack up unanswered.

The AI strategy is a single point of failure, leadership has no one to ask for ROI numbers, and the work she does on this never shows up in her performance review.

03

## Output variance with no measurement

The team can feel that AI tooling is helping, but no one can name where, by how much, or in which workflows. The number in the budget deck ("20-30% faster") came from a one-question survey, not from anything measured.

The next budget review is going to ask a question you can't defend, and "we feel faster" won't survive scrutiny when the CFO asks how you know.

04

## Duplicated tooling spend

Cursor and Copilot are both in production. Roughly \$3,800-5,600/year of overlap. There is no defensible answer to "why both" beyond "we never killed one."

The engineering director will get asked to justify both in the next budget cycle, and the absence of a clear answer hands the budget conversation to someone who isn't in the room.

## YOUR SOLUTION

Solving these is straightforward. No new headcount, no new tools, no multi-quarter program. The approach: substrate first, ownership second, measurement third. Each move sets up the next.

### 01

#### Make convention files load-bearing

Adopt the custom `AGENTS.md` shipped with this assessment. Add a `CODEOWNERS` entry that requires the named owner's review on changes to convention files. Run a 90-minute working session per repo to fill the empty sections from real engineering experience. Set a quarterly refresh ritual.

A new engineer can onboard to "how we use AI here" in 30 minutes by reading the file. That onboarding speed compounds with every hire.

### 02

#### Formalize the AI tooling owner

Write a one-page role addendum (not a full job-description rewrite) that names 8 hours/week of allocation and three quarterly deliverables: convention file updates, ROI report on one tracked workflow, one workflow audit. Announce it in the next engineering all-hands.

A single legitimate first stop for AI tooling questions across all teams. The announcement is the move that makes it real to other teams.

### 03

#### Measure one workflow for 30 days

Pick test scaffolding as the first measured workflow (highest signal-to-noise). Add a one-line tag to the PR template: "AI-assisted? [yes/no]." Pull the data weekly, track for four sprints, and don't try to measure everything at once.

A defensible ROI number for one workflow in 90 days. The number is the entry point to the next maturity level — and the answer to the question leadership keeps asking.

### 04

#### Kill the tool overlap

Run a 30-day usage audit on Cursor + Copilot. Decide which one survives based on actual usage, not preference. Cancel the other before the next renewal. Communicate with 30 days notice, and run a one-hour office hours session for anyone struggling with the transition.

\$3,800–5,600/year saved, defensible answer ready for the next budget meeting, and one less "which tool" question in the AI tooling Slack channel.

---

## WHY ME

A decade inside engineering orgs at Ford and Morningstar. Built and shipped systems that other people now run in production. The IoT telemetry pipeline at Ford cut \$5M in operational costs. A compliance-reporting workflow at Morningstar that was 90% manual now runs as background automation.

I'm building public proof of this consulting practice in real time — these first five AI Readiness Assessments are priced at \$99 specifically so I can publish the case studies. You're seeing a methodology, not a polished testimonial wall, and that's deliberate. The deliverable below is the proof.

### Why this matters to you:

- I've solved versions of these problems before, inside large engineering orgs, with the same constraints you're working with: real budgets, real political pressure, real team politics around tool choice.
- I focus on what ships, not what looks good in a deck. Tight scope, measurable outcomes, no 40-page deck.
- The only metric I care about is whether your team's AI tooling moves engineering velocity in a way leadership can defend. Everything else is secondary.
- I don't treat AI as a trend. Used correctly, it compounds team output. Used poorly, it generates noise that the senior engineers spend their afternoons cleaning up.
- No AI slop. No bloated workflows. The custom `AGENTS.md` shipped with this assessment is meant to be evaluated by your VP Eng in 90 seconds. If it's generic, throw it out and tell me. That's the deal.

## WHAT WORKING WITH ME LOOKS LIKE

STEP	WHAT HAPPENS	WHY IT MATTERS
01	<b>Async intake</b> (~10 min). You fill out a short form covering stack, team size, current AI tools in use, top pain points, and whether you can share repo access.	Lets me show up to the call with context, not start from scratch. The call is for diagnosis, not data collection.
02	<b>45-minute discovery call.</b> We walk through the AI Readiness Quiz live (the same 5 questions on the public article, with probes for the verbatim answer behind each). I capture diagnosed level + 3-4 highest-ROI actions during the call.	The diagnosis happens here. Live answers differ from form answers; the variance is the signal.
03	<b>5 business days</b> , then 3 artifacts arrive in your inbox: this written assessment specific to your team, a Custom AI Adoption Playbook adapted to your stack, and a ready-to-commit <code>AGENTS.md</code> for your main repo.	The deliverable is concrete and yours to keep regardless of what comes next. The <code>AGENTS.md</code> alone usually justifies the fee — VP Eng can evaluate quality in 90 seconds.
04	<b>Optional 30-min follow-up</b> after you've read everything. Free, no commitment. Use it for clarifying questions or to scope deeper engagement if there's an obvious next move.	You're paying for diagnosis and substrate, not a sales pipeline. The follow-up is a service, not a closing call.

## WHAT'S NEXT

Three things, in order, after you read the bundle:

- 1. Drop the AGENTS.md into your main repo as a draft PR.** Get three engineers to review it. The PR conversation surfaces what's actually true about your team's conventions vs. what I inferred. That's where the file becomes load-bearing.
  - 2. Decide which of the 3-4 actions you'll start in the next 30 days.** Don't try to run all of them in parallel. Pick the one with the clearest owner and the lowest political cost.
  - 3. Book the optional follow-up call** if you want a sounding board on sequencing or want to scope a deeper engagement (implementation sprint, retainer, or dedicated workflow build). No pricing here. Those conversations happen on the call once the work is sized to your situation.
-

## THE FAIR-DEAL CLAUSE

If after the 1-hour call I don't see a clear opportunity for your team within reach in 90 days, I refund the fee. Automatic, no email back-and-forth. If you're already at a higher maturity level than your intake suggested and don't need an outside read, I'd rather you keep the money and the hour than walk away from a deliverable that wasn't useful.

---

This is a sample assessment. Your real one will be specific to your team and the conversation we have. Same structure, same voice, different details.

Ready for your own? [Book the AI Readiness Assessment →](#)