FIELD NOTE

October 24, 2026
6 min read

Serious AI starts with the work you keep postponing

AI makes technical debt and undocumented decisions harder to ignore. Why the boring groundwork comes first, and where to start: one process that matters.

IN SHORT

- Serious AI starts with boring work: understand your systems, make the upgrades you postpone and write down how work really gets done.
- On October 21 the PHP project fixed two critical vulnerabilities. Signs of one underlying problem dated from 2024, and fixes were delayed.
- Finding a problem and being able to deal with it are two different capabilities. AI helps with both, but people still decide and own the work.

WRITTEN BY

Mo Kahlain

WITH

Nora Inocta

AI AGENT · FOUNDER'S ASSISTANT

Read online: https://inocta.io/lab/field-notes/serious-ai-starts-with-the-work-you-postpone/

If you want to implement AI seriously in your company, you need to be ready to do the boring work.

That means understanding the systems you depend on and dealing with the upgrades you keep postponing. It also means writing down how work actually gets done, including the decisions that only one person knows how to make.

The recent PHP security vulnerabilities brought this into focus for me.

On October 21, the PHP project released fixes for two critical vulnerabilities. The story involved AI-assisted research, but reducing it to someone asking Claude to find a weakness misses what matters. OpenAI says Calif used GPT-5.6 Sol to identify one flaw, while Wiz independently reported both. The PHP Foundation also describes earlier work with Claude Mythos[1][2].

What happened before the emergency matters more. The Foundation says signs of one underlying problem had appeared in 2024. Fixes involved difficult compatibility choices. Limited resources and a stalled decision process delayed progress. Later reports established the full impact and made the work urgent[1].

That is the part business leaders should pay attention to. Finding a problem and being able to deal with it are two different capabilities. AI can help with both. Between them, people still have to decide and own the work.

01

Technical debt is getting harder to postpone

Technical debt already carried risk before AI. An unsupported component or a fragile integration was never harmless just because nobody had found a way to exploit it.

AI changes how much confidence we should place in that situation staying quiet. The PHP case shows how AI-assisted investigation can expose weaknesses that require a substantial response. I think companies need to reassess the debt they have become comfortable living with.

I would start with a basic question. If a critical update arrived tomorrow, would you know where it applies and who owns the affected systems? Would you know how to test the change?

If the answer requires finding the person who built something five years ago, you have work to do. If nobody wants to touch an application because nobody knows what will break, the delay is already limiting your options.

Not every old system is dangerous, and not every rewrite is worth it. But the cost of leaving things alone deserves an honest review. Keeping a system is a decision. So is continuing to run it without a clear owner or a reliable way to change it.

02

Your systems also need to be usable by AI

Security is one side of the problem. The other is whether AI can do useful work with what you have.

Say an agent helps prepare a customer quote. Pricing sits in one system. Stock is in another. The customer's negotiated terms are in an email. Someone keeps a spreadsheet because the official figures are usually a day behind.

A person may know how to navigate all of that. They know which number to trust and whom to call when something looks wrong. An agent with access to the same screens does not get that judgment for free.

So someone has to decide which information is authoritative and what the agent may change. Then they define how the agent confirms that an action worked. Sometimes a small integration will be enough. Sometimes the process needs simplifying before you connect anything.

Anthropic's engineering guidance makes a practical point here. Tools need clear descriptions and usable interfaces, and they need testing. A capable model still needs a well-designed way to interact with the work[3].

Buying more intelligence while leaving those questions unanswered puts the burden in the wrong place.

03

The business logic cannot stay in people's heads

The same issue exists outside the technology.

A customer gets an exception because someone remembers an old agreement. An order skips a step because an experienced employee recognizes the situation. A manager knows when a delay is acceptable and when it has to be handled now.

That knowledge is valuable. It is also difficult to hand over when nobody has made it explicit.

If you cannot explain how the work happens, what exactly are you asking AI to help with?

A process document that only describes the standard steps will not be enough. The decisions belong in it too. What information changes the answer? Who can approve an exception? When should the work stop? What does a correct result look like?

AI can help interview people and organize what they explain. It can flag where their answers contradict each other. The people responsible for the work still need to check the result and decide which rules should apply. A polished document can be wrong just as easily as a confident answer.

NIST's AI risk framework reinforces that responsibility. It says the roles and responsibilities for managing AI risk should be documented and understood, along with the lines of communication[4]. For a business, that starts with knowing who answers for the result.

04

Start with one process that matters

You do not need to document the entire company or replace every legacy system before starting.

I would pick one recurring process where the friction is real and the outcome can be checked. Preparing a quote or onboarding a customer often works, depending on your business. An internal request can work too.

The useful part is sitting with the people doing it and following an actual case from beginning to end. You see the missing information and the workarounds. You also see the exceptions, which the process everyone wishes they followed never shows.

Then fix what prevents that process from working reliably. It needs an owner and clear decisions, and the information has to be within reach. Only then would I give AI a bounded part of the work and test it against real cases, including the awkward ones.

The measure has to cover the complete result. If drafting takes five minutes instead of an hour but checking and correcting it takes another hour, you have not improved the process yet.

What you learn decides what comes next. This is the thinking behind inocta.io's approach: understand the operation, then decide where AI belongs[5].

05

Give the groundwork a place in the plan

The boring work needs a budget and someone who owns it. Without time on the calendar, it keeps losing to the next demo.

Technical debt and undocumented processes should be part of the AI discussion from the beginning. They affect what you can safely connect and what you can delegate. When something goes wrong, they affect how quickly you can recover.

AI can help with this work too. It can read through code and draft documentation, and it can help test a process. People still have to make the trade-offs and answer for what goes live.

Before adding another agent, I would choose one piece of work and make it clear enough to hand over. That means knowing the systems it depends on and the decisions it requires. And someone has to own the outcome.

That is where I would start.

06

Sources

1. The PHP Foundation, Today's PHP Security Release, and How We're Strengthening PHP Security, October 21, 2026.
2. OpenAI, Patch the Planet: global response coordination as agents uncover critical vulnerabilities, October 21, 2026.
3. Anthropic, Building effective agents, December 19, 2024.
4. NIST, AI Risk Management Framework, Govern 2.1.
5. inocta.io, The Intelligence of Operations.

Read online: https://inocta.io/lab/field-notes/serious-ai-starts-with-the-work-you-postpone/