

MAY 26, 2026

# The knowledge in your head is your AI bottleneck

Senior operators struggle with AI agents because their expertise has compressed into instinct. The agent needs the explanation. The expert can't access it.

By Dave Taylor



AI agents don't work for senior operators because the people who'd benefit most from delegation are the worst at it – not from lack of effort, but because their expertise has compressed into instinct. The senior PM who *just knows*. The strong salesperson who *just feels* it. They can't explain what they do because they no longer think about it consciously. The agent needs the explanation. The expert can't access it. [Berkeley's California Management Review](#) calls tacit knowledge "your next competitive moat" – and the people sitting on the most of it are the ones running headfirst into the wall when they try to delegate to AI. The cruel irony: the more experienced you are, the harder it is.

## Why experts can't explain what they do

When you're new at something, everything is deliberate. You follow checklists. You think through each step. You can tell someone exactly what you're doing and why because you're still figuring it out yourself. That changes with experience. Think about driving a car. After 20 years, you don't think about turning right. It just happens. Your brain automated the process so it could free up capacity for harder things — like deciding whether to take the scenic route or watching that cyclist who's about to do something unpredictable. Knowledge work follows the same pattern. Early in your career, you think through every decision. As you get better, those decisions compress into reflexes and gut feelings. You stop reasoning through decisions. You just make them.

The thing that makes you fast and effective is the same thing that makes your knowledge invisible. It's been compiled from something you could explain into something you just do — and you no longer have the explainable version. Take a really experienced product manager. They don't consciously think *"I should cross-reference the revenue dashboard with the churn data before forming an opinion."* They open three tabs, glance at the numbers, and just know. If you asked them to explain what they just did, they'd narrate backwards from the conclusion. They'd give you a clean story. But it wouldn't be the real process — the real process involved dozens of micro-evaluations drawing on years of pattern recognition across multiple companies and products. That pattern matching is what makes them valuable. It's also what makes their work nearly impossible to hand to an AI agent.

The agent doesn't want a clean narrative. It wants the messy, specific, step-by-step version the expert stopped using years ago. Same goes for a strong salesperson — they don't consciously decide to mirror a prospect's phrasing or slow their cadence when they pick up on defensiveness. They just do it. A senior engineer doesn't explicitly reason through a concurrency issue; they feel it, say *"that's bad"*, then go look and they're right. Everyone calls it experience. This isn't a failure of self-awareness. It's the intended outcome of getting good at something. Expertise is supposed to become automatic. That's what makes it powerful. But it's also what makes it a nightmare to delegate — to humans or to machines.

*The thing that makes you fast and effective at your work is the same thing that makes your knowledge invisible — to yourself and to any agent you try to delegate to.*

## Three problems this has always caused

Tacit knowledge isn't new and the problems it causes aren't new either. Agents just made them urgent.

**Delegation fails.** Every management book talks about this. The standard explanation is that managers are control freaks. The real explanation is simpler: they don't know how to express what's in their heads. It's not that they won't let go. It's that they can't describe what they're letting go of in enough detail for someone else to do it well. So they either micromanage or they hand it off and get disappointed.

**Good people don't get promoted.** The most common reason strong individual contributors plateau is that they can't be replaced. Their knowledge is locked in their head. Even when they're qualified for the next level, nobody wants to risk losing their expertise in the current role. The very thing that makes them great becomes the trap that keeps them there.

**Knowledge walks out the door.** When someone leaves, years of institutional knowledge goes with them. Not the documented stuff – the real stuff. The unwritten rules, the relationship context, the *we tried that in 2019 and here's why it didn't work* stories that save teams from repeating mistakes.

Organisations have spent decades trying to fix these problems with top-down knowledge management initiatives, wiki projects, documentation sprints. None of it worked at scale because there was never a direct personal incentive. The benefit went to the organisation. The person who documented their expertise was often the person who made themselves replaceable. So the work didn't happen.

## Why agents change the incentive

Here's what's different now: agents flip the incentive structure on its head. For the first time, the person who puts in the work to articulate their expertise is the person who directly benefits. Your agent gets better. Your work gets faster. You get compounding returns because the agent is learning from a rich, detailed understanding of how you actually operate. You're not writing a wiki that nobody reads. You're not filling in a knowledge base so HR can tick a box. You're building the operating system for your own personal leverage tool.

And once you've done it once, the second time is faster. The tenth time takes minutes. Because you've built the muscle of understanding and expressing what you actually do – not the job description version, the real version. This is a bottoms-up knowledge

management revolution dressed up as a consumer AI product. Most people — including the people building these products — don't see it yet. If they did, they'd invest a lot more in the onboarding experience than three setup questions and a *you're good to go* button. The selfish incentive is the unlock. The person who documents their expertise is the person who gets the leverage.

## The uncomfortable divide

This leads somewhere uncomfortable. If the value of agents depends on your ability to articulate your work, agents are about to create a very visible divide. Right now, tacit knowledge is invisible. Nobody knows you can't describe your own process because nobody's ever asked you to. Performance reviews measure outputs, not self-knowledge. You can be a phenomenal operator with zero ability to explain how you operate, and the system has never penalised you.

Agents are about to change that. [Microsoft's 2026 Work Trend Index](#) shows the gap forming already: among advanced AI users, 80% say AI lets them spend more time on high-value work — compared to roughly two-thirds across all users. The advanced users aren't using a different model or a different platform. They've put the work into articulating their expertise. Their agents get better because their context gets better. In a world where everyone has access to the same tools, the differentiator isn't which model you use or how much you spend on infrastructure. It's whether you can feed the thing well enough to get results.

The people who skip that step install the agent, play with it for a weekend, hit the wall, and walk away thinking agents don't work. They'll be wrong — the agent worked fine. The problem was never the technology.

## What to do about it

If you're reading this and thinking *that sounds like me* — good. That's the first step. The fix isn't to sit down and try to write everything you know in one marathon session. That almost never works, because the knowledge you need to capture is exactly the knowledge you can't access by just thinking about it. The 40-hour-and-it-still-didn't-work pattern we covered in [Part 1 of this series](#) is what happens when you try to brute-force the extraction without structure.

What works is **structured extraction**. Having someone — or something — ask you the right questions in the right order, with the right follow-ups, so the knowledge surfaces naturally. Not *"what's your job?"* but *"walk me through what you actually did last Tuesday."* Not *"how*

*do you make decisions?" but "tell me about the last hard call you made and what you were weighing up." The questions that surface tacit knowledge are always specific, always grounded in real recent work, and always followed by "what made that obvious to you?"*

That's the work we do at CxD. We help leaders and teams extract the operational knowledge that's been compressed into instinct, turn it into structured context, and use it to deploy agents that actually work. Not because we're selling you an agent – you've already got one. We're helping you get value from it. Start with one question this week: *if I had to hand my job to someone tomorrow, what are the ten things I'd need to explain that aren't written down anywhere?* Write them down. That list is your starting point.

If you'd rather skip the cold-start and run the extraction with someone who's done it before, [book a 30-min working session](#). We'll find the slice worth surfacing first.

*Related: this is Part 2 of the AI Agent Reality Check series. [Part 1 – Treat your AI agent like a new hire](#) covers the onboarding mechanics. Part 3 (failure modes) follows next week. And [our case studies](#) for what shipped AI work actually looks like.*