

FIELD GUIDE 08

Applied ideas from
Mike Taylor's Eight Levels
of AI Adoption (Every)

THE IGNITE ALMANACK OF
**THE AI
LADDER**

Eight rungs, two questions, and the one move worth making this month.

A FIELD GUIDE FOR UK SMES

IGNITEAI SOLUTIONS

THE PROBLEM

One viral post can make you feel like you're doing AI wrong.

Someone's agent answers their email while they sleep. Someone else runs twelve AI sessions in parallel. Meanwhile you're pasting things into ChatGPT and wondering if that still counts.

It counts. The mistake is thinking AI adoption is a race to the top. I sit with MDs every week who feel behind, and when we map what they actually do, the problem is never the level they're at. The problem is that every task in the business sits at the same level. The payroll and the meeting notes get the same treatment, which means one of them is over-supervised and the other is a risk nobody has looked at.

AI use is not a competition.

Mike Taylor, Every

Taylor's eight-level map is the clearest picture yet of how far you can hand work to AI. This guide is the owner's version of it: what each rung looks like inside a UK SME, which rung each task belongs on, and how to climb without hiring an engineer.

WHY THIS MATTERS

The firms one rung up from you are not smarter and they did not spend more. They stopped asking "how good is AI" and started asking "how much of this task can it own". Different question, different results.

THE IDEA

Eight rungs. Every task on its own rung.

Mike Taylor, who runs tech consulting at Every, mapped eight levels of AI adoption, from asking a chatbot a question to a manager agent running a team of agents. Here is the ladder translated into plain business terms.



The eight levels, owner's translation. Framework: Mike Taylor, "The Eight Levels of AI Adoption", Every, 2026.

The rung a task belongs on comes down to two questions. How much do you trust the output on this task? What does it cost you if it's wrong? High stakes and low trust: stay low, keep approval in the loop. Low stakes and proven trust: push it up and stop babysitting it.

WHERE IT LANDS

The same business, eight tasks, five different rungs.

Run the two questions over a normal week and the ladder sorts itself. This is what it looks like inside a typical 50 to 200 person firm.

The task	The rung	Why it sits there
Payroll, VAT, anything HMRC sees	3	High stakes. It prepares, a human approves every step.
Complaint replies	3	Your reputation in every word. Draft fast, approve everything.
Client proposals	3 to 4	Your name on it. It drafts from your files, you shape and sign off.
Chasing overdue invoices	4	Low risk, proven pattern. It drafts the chase, you skim before send.
Market and competitor research	4	A wrong turn costs minutes. Check the result, not the process.
Meeting notes and actions	6	Runs after every call. Nobody types up minutes again.
The morning numbers: cash, pipeline, diary	6	Assembles itself before 8am. You read it with your brew.
A team of agents running a function	7 to 8	Real, and mostly not worth your money yet. Revisit in a year.

THE PATTERN UNDERNEATH

Nothing in this table needed new software choices or an IT project. Same tools, different rung per task. The decision is the work.

The top half of the ladder does not need engineers on payroll.

Taylor is straight about his map: he places knowledge workers at Levels 1 to 4 and calls the top half engineer territory. For big-company software teams, fair. For an SME, the picture we see on the ground is different, and it matters for what you buy next.

In our own business the morning brief writes itself before anyone has poured a brew. The weekly pipeline triage runs on a schedule. Meeting transcripts file themselves into a knowledge base overnight. That is Level 5 and 6 work, running daily, and nobody here writes code for a living.

The honest part: someone does have to build the scaffolding once. The playbooks, the guardrails, the quality checks, the rule that says a human signs off anything a client sees. That is a few weeks of setup, done once. Running it afterwards is a management habit, the same discipline you'd apply to a new starter. You wouldn't hand a squaddie a rifle and skip the range package. Same rule here.

So the question for an owner is never "are we ready for agents". It is "which one task earns a rung up this month, and who owns it when it moves". Most firms haven't asked it yet.

92%

of UK directors see AI potential (IoD, 2025)

16%

feel prepared to act on it (IoD, 2025)

43%

of small firms have no AI plans at all (FSB, 2025)

HOW TO START

One task. One rung. Two weeks.

- 1 List the drudgery.** Ten tasks that drive you or your team crazy every week. Real ones, from the diary and the inbox, in ten minutes.
- 2 Score each task twice.** Trust: how good is AI at this today, high, medium or low? Stakes: what does a mistake cost, high, medium or low?
- 3 Place them on the ladder.** High stakes stay on rungs 1 to 3 with approval built in. Low stakes with high trust are your climbers.
- 4 Pick one climber. Just one.** Write down what good looks like before you start. If you can't describe how the task is done today, it isn't ready to hand over.
- 5 Move it one rung and give it two weeks.** Count the minutes it gives back. Keep the human checkpoint until the output has earned its way past you.
- 6 Name an owner.** Every task above rung 4 needs one person who keeps its instructions current. Systems die of neglect, and nobody notices until the output goes stale.

TRY THIS ON MONDAY

Take the ten-task list into Monday's management meeting and **agree the one task that moves one rung this month.** One decision, fifteen minutes, and the ladder stops being a LinkedIn post and starts being your operating model.

Three prompts that do the sorting for you.

Paste these into Claude or ChatGPT as they are. Fill in the brackets. Each one takes minutes and gives you something concrete to act on.

PROMPT 1: LEVEL-CHECK MY WEEK

Here are ten recurring tasks in my business: [paste your list]. For each task, score AI trust (high/medium/low) and stakes if it goes wrong (high/medium/low). Place each on an 8-level ladder from "chatbot" to "runs without being asked". Tell me which single task is the best candidate to move one level up this month, and why.

Turns the ladder into a sorted to-do list for your actual week.

PROMPT 2: ONE RUNG UP, SAFELY

I want to move [task] from "I do it with AI help" to "AI does it and I review". Design a two-week pilot: what good output looks like, the checks the AI must run on its own work, where a human must approve, and the one number that tells me it worked.

Builds the guardrails before the climb, so the pilot can't quietly fail.

PROMPT 3: THE 7AM TEST

If a morning briefing assembled itself for my business at 7am, what should be in it? I run a [type] business with [n] staff. My week revolves around [cash/jobs/pipeline/rota]. List what the briefing would need to read, what it would flag, and the three decisions it should tee up for me by 8am.

Shows you what rung 6 would feel like before you spend a penny on it.

THE NEXT STEP

CLIMB ONE RUNG.

IgniteAI builds the scaffolding that lets a normal business run the top half of the ladder: the playbooks, the guardrails, the daily briefings, the systems your team runs without us. We start with the friction, never the technology.

Skills that stay when we go.

If the ten-task list surfaced something bigger, that conversation costs nothing.

IGNITEAI SOLUTIONS

igniteaisolutions.co.uk · chris.duffy@igniteaisolutions.co.uk

THE IGNITE ALMANACK
FIELD GUIDE 08