Asking the Questions

Adopting deliberately — one narrow job, real rules, and a business that gets better at learning

Parts One and Two gave you the vocabulary and the judgement. This final guide is about method: the specific, unglamorous practice that separates the businesses seeing measurable returns from the majority holding subscriptions.

The reframe for this territory: nobody owns the answers here — not the vendors, not the consultants, and not the people posting confidently online. The field is moving too fast for expertise to settle. What you can own is the questions, and a method for learning faster than the businesses around you. In a domain changing this quickly, that is a more durable advantage than any tool.

Everything here is doable this quarter, without a consultant, a budget line, or a technical hire.

John Obidipe · catalystgrowthcoach.co.uk

01

The adoption method — one narrow job

Pick the job, not the tool

The idea: Start from the hours map in Part Two: one specific, repetitive, high-volume task where being wrong is cheap and recoverable. Then find the tool that does it — rather than buying a tool and hunting for uses, which is how the dabbler's plateau begins.

How to use it: Narrow beats ambitious every time at this stage, because narrow is measurable. 'We use AI for quotes' can be tested in a month; 'we're adopting AI' cannot be tested at all.

The question to ask: *"What is the one task where this would save the most hours with the least risk?"*

Name an owner and a review date

The idea: One person accountable for making it work, and a date — six or eight weeks out — when you decide together: adopt properly, adjust, or stop.

How to use it: This single discipline prevents two of Part Two's five patterns outright. And the decision must be real: a pilot that cannot be stopped is not a pilot, it is a purchase with extra steps.

The question to ask: *"Who owns this, and when do we decide?"*

Measure two things, not twenty

The idea: Hours saved on the task, and how often the output needed correcting. Rough numbers from the person doing the work are entirely sufficient.

How to use it: This is the whole measurement system, and it settles the argument honestly in both directions — including telling you when something impressive isn't actually saving anything once checking is counted.

The question to ask: *"How long did this take before, how long now, and how often do we have to fix it?"*

Expand from evidence, then write it down

The idea: When a job works, capture how: the prompt that produces good output, the checking step, the exceptions. Then move to the next job on the map.

How to use it: This is where compounding happens. The prompts and rules that work are genuine business assets — they belong in a shared document, not in one person's head or chat history. Businesses that write them down get better at AI monthly; businesses that don't restart every time someone leaves.

02

The rules your business needs — all four of them

Not a policy document. Four sentences, said out loud to everyone, written on one page:

- **What never goes in.** Client data, payroll, contracts, anything confidential — unless it is the approved tool with the right protections. Name the boundary and name the sanctioned alternative, or people will simply route around you.
- **What always gets checked.** Anything factual, numerical, legal, or reaching a customer. Draft freely; verify before it leaves the building.
- **Never take agreement as evidence.** These tools tend to side with whoever is asking. When you use one to think rather than to draft, ask it for the case against your plan, not whether your plan is good. A rule worth saying out loud, because it is the failure nobody notices.
- **Who is accountable for the output.** The person who sends it, always. 'The AI wrote it' is not a defence and should never be an acceptable sentence in your business.

- **When we tell people.** Your position on disclosing AI use to clients — which for most businesses is: freely used for drafting and analysis, always reviewed by a person, never pretending to be a human relationship it isn't. Decide it before a client asks, not during.

03

Agents — autonomy, earned

The 2026 shift from AI that answers to AI that acts is real, and it will reach your business through the software you already use. The framework for staying deliberate:

- **Start where errors are cheap and reversible.** Internal drafting, sorting, monitoring, preparing. Not payments, not contracts, not first contact with a customer.
- **Keep a human checkpoint until evidence says otherwise.** The agent proposes; a person approves. Remove the checkpoint task by task, only after watching it be right consistently — and only where being wrong is survivable.
- **Ask for the action list, and shorten it.** "What can it do without asking?" Any competent provider can answer precisely. Begin with the shortest version of that list you can live with; expand deliberately.
- **Watch for silent failure.** The dangerous agent errors are not dramatic — they are quiet, repeated, and undetected: the follow-up never sent, the record wrongly updated, the enquiry misrouted for a month. Whatever you automate, keep a way of noticing nothing happened.
- **Check it is actually an agent.** 'Agent' is 2026's most inflated word, and plenty of what carries the label is a scripted automation with a new price. A real agent chooses its next step. Ask what decides that, the model or a rule somebody wrote, because the answer changes both what it can do for you and how carefully you need to watch it.

04

Your people — the question nobody asks

The question that separates a leader from a buyer of software: what is this doing to the capability of my team?

AI genuinely raises the floor — people produce better first drafts, structure their thinking faster, get further alone. It also, used carelessly, lets junior people skip the struggle that builds judgement, and lets experienced people stop practising the craft that made them worth listening to. Both effects are real and both are already happening in businesses like yours.

The practical stance: be explicit that AI drafts and humans judge, insist that people can still do the core thing without it, and keep the tasks where skill is genuinely built as human work for the people building it. Then say so openly — a team that knows why the line is drawn where it is will hold it when you're not watching.

05

In the room — the twelve questions

The distillation of the series. Any of these, asked plainly, makes you the most useful non-technical person in an AI conversation:

1. "Show me that working on our stuff, not the demo."
2. "Where do the errors go — who checks, before it reaches anyone?"
3. "What happens to our data — is it used for training, and who can see it?"
4. "What changes on a normal Tuesday if we do this?"
5. "Is this augmentation or automation — and are we attempting the second too early?"
6. "What can it do without asking a human, and can that list start shorter?"

7. “Who owns this, and when do we decide to adopt, adjust, or stop?”
8. “How long did this take before, how long now, and how often do we fix it?”
9. “What can it do now that it couldn’t two years ago?” (for anything sold as new)
10. “What is this doing to what my people can do without it?”
11. “Is this genuinely an agent, or an automation being sold as one at agent prices?”
12. “Is it agreeing with me because it’s right, or because it agrees? Give me the case against.”

The end of feeling behind

Three guides ago, AI was a wall of noise you were supposedly already losing to. Now: you have the words (Part One), you can read the pitch and locate the real value (Part Two), and you have a method — one narrow job, an owner, two measures, evidence before expansion (Part Three).

You are not an AI expert, and in this territory that matters less than anywhere else, because there are fewer real experts than the volume of confident opinion suggests. What you have is more durable: the questions, and a business that learns deliberately while others accumulate subscriptions.

And the closing observation this library keeps arriving at: what looks like an AI problem is usually a structure problem in new clothing. The dabbler’s plateau is an ownership gap. The abandoned pilot is a decision nobody made. The shadow tools are a signal about friction your people couldn’t get fixed. The best businesses will not be the ones that adopted AI first. They will be the ones that were already clear about how decisions get made, who owns what, and how they learn — because those are the businesses where a powerful new tool compounds instead of scattering. That clarity is the work underneath the work, and it is where the conversation goes next.

John Obidipe is a business transformation specialist and Gallup-Certified CliftonStrengths coach. He works with founders and leadership teams on the structural side of growth — the part the numbers can only hint at. · coach@catalystgrowthcoach.co.uk · catalystgrowthcoach.co.uk