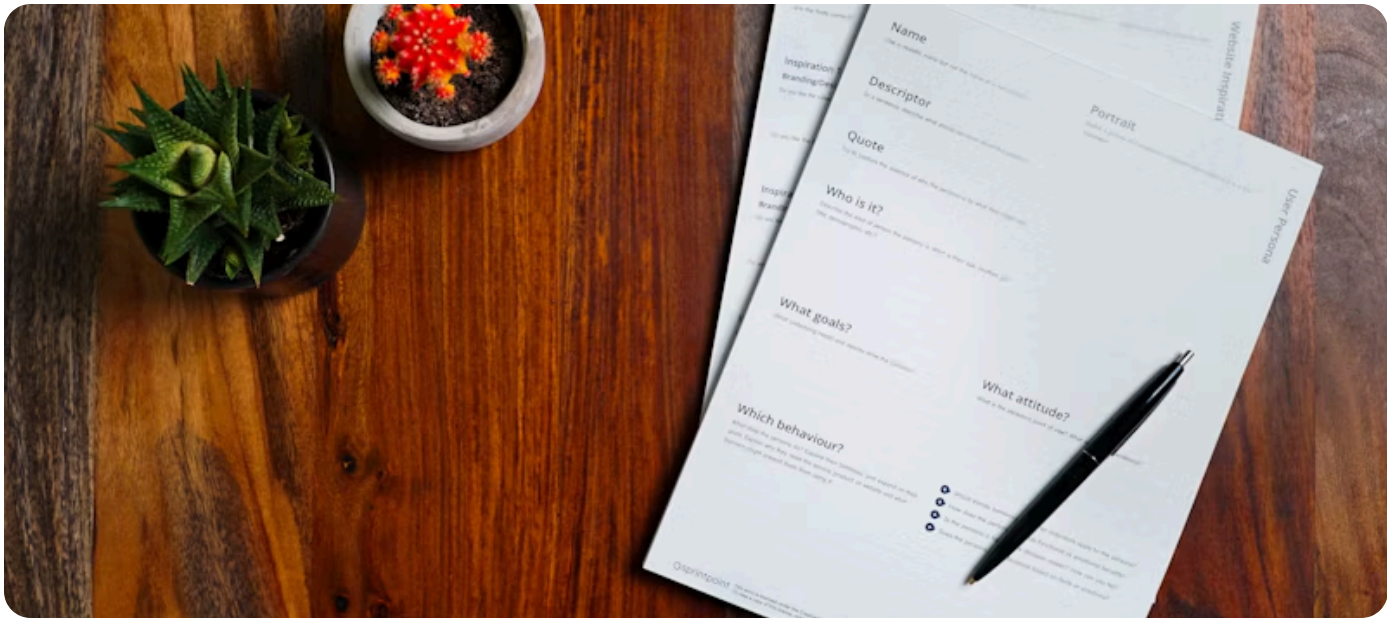
Article

You've tried ChatGPT. That's not the whole story.

AI in plain English: what it takes to trust AI with real work, in four industries, with the words you will hear and a first step.



Bo Thompson, Founder & CEO · November 10, 2026 · 9 min read



Starting is easy.

Anyone can open an AI assistant today and get real work done with it.

Making it work takes the right setup.

Your data kept safe, answers you can check, a person approving what matters, and someone who has done it before.

A story you may recognize



Every Friday, the office manager at Ferncliff Contracting, a 45-person builder, chased timesheets. Her first try at AI was a chatbot. It worked for a week. Then it got a pay rate wrong, and nobody caught it.

The crew foreman stopped trusting it, and that was the end of it.

The fix wasn't a better tool. It was the setup around the tool: mapping how timesheets really flow, deciding where a person checks, and locking down what payroll data the AI can see. That setup is what this page is about. (Illustrative example.)

AI on your own vs AI set up for your business

Today’s assistants are genuinely powerful. Many can read your email, open your files and act in your apps. That’s exactly why setup matters.



ChatGPT, Claude and similar tools can do the work. The question is whether you can trust them with it. Used on your own, you get the capability. Set up for your business, you also get security, accuracy and accountability.

	AI on your own	AI set up for your business
What it's allowed to touch	Whatever the user connects, often with full access to their email and files.	Only what the task needs, set on purpose and reviewed regularly.
Getting the facts right	Answers from general knowledge. Usually good, sometimes confidently wrong.	Answers from your own approved records, with the source shown. If the record doesn't say, it's set up to say so instead of guessing.
Checks before it acts	Depends on the user remembering to check.	Approval steps built in for money, contracts, customers and safety.
Consistency	Varies by who asks and how they ask.	The same process every time, tested before any change goes live.
Record and accountability	Scattered across personal chats.	A log of what it did and why, and a named owner who answers for it.
When something goes wrong	You may not notice.	Monitoring that flags it, an off switch, and a way to roll back.

CRM means your customer database. ERP is the system that runs orders, inventory and billing. Both come up below.

What AI can actually do, in your world

The best first jobs for AI are the boring ones: work your people repeat every week, that follows a pattern, and that someone can check. Here is what that looks like in four industries we work with.

Each story is an illustrative example. The companies are made up and the hours are estimates. Every one shows what the AI does, where a person checks, and the setup that made it work.

Ferncliff Contracting: timesheets from voice notes

Construction, illustrative

Before

After

Every Friday the office manager chased crews for timesheets, decoded photos of paper, and retyped it all into payroll.

About 4 hours a week (estimate)

Crews leave a voice note from the truck. The timesheets are waiting for her to approve.

Minutes of review instead (estimate)

What the AI does

Turns each voice note into hours by person and job, matches them to real job codes, and flags anything odd, like a missing day or a 14-hour shift.

Where a person checks

The office manager approves every timesheet before payroll. The AI never changes a pay rate.

The setup that made it work

Mapping how timesheets really flow, connecting only the crew list and job codes, and keeping pay data out of its reach.

Prairie Ridge Property Co.: maintenance requests sorted and routed

Real estate, illustrative

Before

Requests arrived by email, text, portal and voicemail. A coordinator read each one, guessed the urgency and passed it on.

About 6 hours a week (estimate)

After

Every request is sorted by urgency and trade, a work order is drafted, and the tenant gets a reply that someone has it.

About 1 hour a week of checking (estimate)

What the AI does

Reads the request, tags it (plumbing, HVAC, appliance), sets urgency using the company's own rules, and drafts the work order and the tenant reply.

Where a person checks

Anything about water, gas, power or safety goes straight to a person. The coordinator approves each vendor dispatch.

The setup that made it work

Writing the urgency rules with the maintenance lead, and connecting only the work-order system.

Brasswind Industrial Sales: follow-up after every call

Sales teams, illustrative

Before

After

Reps meant to follow up after calls. Many didn't, and CRM notes were thin or missing.

About 3 hours a week per rep on notes and emails (estimate)

Within minutes of a call, a follow-up email and CRM update are drafted and waiting for the rep.

A few minutes per call to review (estimate)

What the AI does

Takes the rep's call notes, drafts a follow-up in the company's voice, fills in the CRM fields, and sets the next-step task.

Where a person checks

The rep reads and sends every email. Nothing goes to a customer on its own, and it never promises pricing.

The setup that made it work

Templates written with the sales manager, CRM access limited to each rep's own accounts, and testing on past calls first.

Kettlebrook Supply Co.: quote drafts

Distribution, illustrative

Before

Inside sales built quotes by hand from emailed requests, looking up part numbers, price and stock in the ERP (the system that runs orders, inventory and billing) one line at a time.

About 10 hours a week across the team (estimate)

After

The quote is drafted before the rep opens the email. The rep checks it, adjusts and sends.

About 2 hours a week of review (estimate)

What the AI does

Reads the request, matches the customer's part names to Kettlebrook's catalog, pulls current price and stock, and drafts the quote.

Where a person checks

A rep reviews every quote before it goes out. Any discount off list goes to a manager.

The setup that made it work

A read-only connection to the ERP, a cleaned-up part-number cross-reference, and a test run on last month's quotes.

The words you’ll hear

You don’t need to be technical to lead an AI project, but the vocabulary helps. Here are ten terms you’ll hear in any vendor meeting, in plain English.

AI model (large language model)

The engine underneath tools like ChatGPT. It has read a huge amount of text and predicts good answers. Smart, but it knows nothing about your company until you connect it.

Prompt

The instructions you give the AI. Better instructions get better results, but a great prompt can't fix missing data or a broken process.

AI agent

An AI that works toward a goal in steps: it reads, decides, uses your software and checks results, inside limits you set.

Integration

Connecting AI to the systems you already use (your CRM, ERP, accounting or job software) so it works with real data.

Hallucination

A confident wrong answer. The AI fills a gap with something plausible. The main fix is grounding, backed by checks and approvals.

Grounding

Making AI answer from your approved records instead of memory, and show where the answer came from.

Guardrails

The rules and limits around the AI: what it can see, what it can do, and what it must never do.

Human in the loop (approval step)

A point where a person reviews and approves before anything important happens, like a payment, a contract or a customer email.

Governance

Who decides which AI tools are used, what they may touch, and who answers for them. The rules that let you say yes safely.

MCP

A common plug that lets AI connect to business software. Short for Model Context Protocol.

Why a proper implementation can be trusted

AI earns trust the same way a new hire does: clear instructions, limited access at first, someone checking the work, and a record of what was done. The model matters less than the work around it.

Remember Ferncliff's chatbot from the top of this page. The tool wasn't the problem; the missing setup was. Plenty of companies are where Ferncliff started. In the Fed's Small Business Credit Survey, 2026 Report on Employer Firms, 46% of firms use AI; 7% of AI users have fully integrated it (source). The steps below are what closes that gap.

Accuracy is part of trust. A proper setup answers from your own approved records, shows its source, and says "I don't know" when the record doesn't cover it. Then it's tested on real past work before it goes live. Nothing is perfect, but you get far fewer wrong answers, and the ones that slip through get caught.

What makes it trustworthy. Start at 1; each step builds on the one before.

1. **Map how the work really flows.** Not the process on paper. The one your people actually follow, with its exceptions.
2. **Connect only what's needed.** The job list, not the whole accounting system. Its answers come from these approved sources, with the source shown.
3. **Limit what it can touch.** It can read prices but not change them; draft a payment but not send it.
4. **A person approves money, contracts, customers and safety.** Built into the system, so it can't be skipped.
5. **Test before it goes live.** Run it on last month's real work and compare, including whether its answers match your records.

6 **Watch it and keep a log.** Every step recorded: what it did, when, and why.

7 **An undo button.** Any change can be rolled back, and the system can be paused.

8 **A named owner.** One person who answers for it, reviews it and decides when it changes.

Security and governance, without the jargon

Good security and governance are what let you say yes to AI. They aren’t red tape: they’re the short list of decisions that make it safe to put AI into real work. Most of them fit on one page.

What never goes into AI tools

Customer personal details, confidential client files, passwords and bank details. Write it down so nobody has to guess.

Where your data goes

Know which vendor holds it, where it's stored, how long it's kept, whether it's used to train their AI, and who on your team can see it.

Answers you can check

The AI answers from approved company sources, shows where each answer came from, and is set up to say so when it isn't sure. That's how you keep most wrong answers out of real work, and catch the rest.

A log of what the AI did

Every action recorded, so you can answer "what happened?" in minutes, not days.

An off switch

Anyone responsible can pause the AI at once, and it falls back to the way the work was done before.

> **Written guidance, in three parts**
>
> **Approved tools:** which AI tools staff can use for work, on company accounts.
>
> **What never goes in:** the short list above, in plain words.
>
> **Who approves new tools:** one named person or small group says yes or no, so a new tool gets checked before your data goes into it.
>
> One page is enough to start. Call it guidance; it's there to help people use AI well, not to scare them off.

For the deeper version (logins, permissions, testing and vendor checks), see the security and governance section of the Agent Stack Field Guide.

The three worries, answered

We hear the same three worries from almost every leadership team. Each one is real, and each one is solvable.

"AI doesn't work"

The tools work. What usually failed was the setup: no approved sources to answer from, no checks and no owner. So the first wrong answer ended the experiment.

What fixes it: one well-chosen task, grounded in your own records, with a person approving the output.

"AI isn't secure"

Any AI tool on personal accounts with no rules is a risk. AI on company accounts, with limited access, a log and clear guidance, is a different thing entirely.

What fixes it: the guidance, the limits and the log, set up before the first task goes live.

"IT won't let us use it"

IT is usually saying no to the unknown, not to AI. Give them a named tool, a clear data boundary and an off switch, and they become partners in the project.

What fixes it: bring IT in at the start, with a plan they can review and approve.

Your first step

Pick one painful task: the kind of work your people repeat every week, that follows a pattern, and that someone can check. Timesheets, maintenance requests, quotes and call follow-up are all good examples.

Then bring it to people who have done this before. We'll look at how that work really flows, what it would need to connect to, and where a person should stay in charge.

See where you stand: score your business in 5 minutes

Go deeper: the Agent Stack Field Guide

Sources and notes

Federal Reserve Small Business Credit Survey, 2026 Report on Employer Firms: 46% of firms use AI; 7% of AI users have fully integrated it.

Ferncliff Contracting, Prairie Ridge Property Co., Kettlebrook Supply Co. and Brasswind Industrial Sales are illustrative examples. Hours are estimates.
