



PineyWoods

AI STRATEGY

The AI Starting Map

Where AI actually earns its keep in a small business - and where it doesn't.

The AI Starting Map

Where AI actually earns its keep in a small business - and where it doesn't.

You don't need an AI strategy. You need a first move.

By now you've heard the pitch from every direction. A vendor emails you twice a week. A competitor put "AI-powered" on their homepage. Someone at a conference said that if you're not using AI yet, you're already behind. The message is always the same - AI is going to change your business - and it always stops at the exact moment you need it most. Nobody hands you the next sentence. Nobody tells you where, on your actual Tuesday, to begin.

So the question sits there, too big to answer. "What's our AI strategy?" is the kind of question that wants a six-month engagement and a slide deck, and you have a business to run. Most owners we meet have taken a run at it once or twice, gotten overwhelmed, and quietly filed it under "someday." That's a reasonable response to a badly framed question.

Here's the reframe that makes it answerable. You don't need a strategy for all of AI. You need to find the one or two places in your business where it would genuinely earn its keep - and start there. Not the most impressive thing. The most useful thing. The real question isn't "should we use AI." It's "where would it pay off first."

This paper is that map. It won't make you an expert, and it won't try. It will give you a calm, practical way to look at your own business and spot the handful of tasks worth starting with - and, just as important, the ones worth leaving alone for now. You can work through it in an afternoon. By the end you'll have a short list, a ranked top candidate, and a way to tell within two weeks whether you picked right.

The "earns its keep" test

Before you weigh one opportunity against another, you need a way to tell a real opportunity from a shiny one. The shiny ones are easy to spot once you know the tell: they demo beautifully and then never touch your real work again. The real ones are usually dull. They're the tasks nobody brags about.

Here's the test we use. Take any task you're considering and ask four questions. A good first project answers yes to all four. Miss even one and it's probably not your starting point - not never, just not first.

Does it happen often? A task you do forty times a week is worth improving. A task you do twice a year isn't, even if it's painful each time. Frequency is what turns "saved a few minutes" into hours back every week. Rare problems are interesting. Frequent problems are where the money is.

Does it follow the same shape every time? This is the one people skip, and it matters most. AI is dependable when a task has a predictable structure - the same kind of input coming in, the same kind of output going out, a clear sense of what "good" looks like. "Sort these incoming emails into three buckets" has a shape. "Use your best judgment on whatever lands today" does not. The more a task looks the same each time, the more you can trust the result.

Does someone already dislike doing it? Pick work your team won't miss. Tedious, repetitive, a little soul-deadening - that's a feature here, not a problem. People forgive a new system quickly when it takes a chore off their plate, and they fight it when it touches work they're proud of. The groan when a task comes up is a good sign you've found the right one.

Is the cost of a wrong answer low enough to start here? Everything makes mistakes - people, software, AI. The question is what a mistake costs and how easily you'd catch it. A first-draft email a person reviews before sending is a safe place to be wrong; you see the slip and fix it in seconds. Wiring money based on an unchecked AI reading of an invoice is not. Start where errors are cheap and visible. You can move to higher-stakes work later, once you trust the thing.

Four questions: often, same, disliked, safe to be wrong. Hold a task up to the light and ask all four. The ones that answer yes across the board are your candidates. Write them down - you'll rank them next.

Effort vs. impact: ranking what's left

The test gives you a short list. Now you need to choose among them, because you should start with exactly one. The simplest way we know is to weigh each candidate on two things: how much it would help if it worked (impact), and how much work it takes to get there (effort).

You're looking for the boring sweet spot - high impact, low effort. Not the moonshot that would be amazing and takes six months. Not the easy thing that barely moves anything. The task that helps a lot and is quick to stand up. Those exist in almost every business, and they're usually hiding in plain sight because they're unglamorous.

Effort, for a small business, is mostly about three things: how clean and reachable the data is, how many tools have to talk to each other, and how high the stakes are if it's wrong - higher stakes mean more checking, more setup, more care. Impact is simpler: how much time or money it frees up, and how often.

Let me make it concrete. Say you run a 25-person property management company. You run three candidates through the test and all three pass, so now you rank them.

The first is drafting replies to routine tenant emails - the same maintenance questions, lease questions, and "when is rent due" notes, over and over. In our experience a team like this fields

something like fifty of these a day, a few minutes each. The shape is consistent, everyone dislikes it, and a person reviews every draft before it sends, so a wrong answer costs almost nothing. High impact, low effort.

The second is summarizing long vendor contracts before renewal. Genuinely useful - but it happens maybe a dozen times a year. Helpful when it happens, not frequent enough to move your week. Medium impact, low effort.

The third is predicting which tenants are likely to leave so you can get ahead of vacancies. It sounds like the exciting one. It's also the one that needs clean historical data you may not have, judgment that's hard to pin down, and a real cost if it's wrong - you might hassle a happy tenant or miss an unhappy one. High effort, uncertain impact, higher cost of a mistake.

Ranked honestly, the email drafting wins, and it isn't close. It's the least impressive of the three and the one most likely to quietly give the office an hour or two back every day. The contract summaries are a fine second project once the first one is paying off. The churn prediction goes on a "maybe later" list - not because it's a bad idea, but because it's a bad first idea.

That's the whole exercise. Rank by impact over effort, be honest about the effort, and start with the unglamorous winner.

Where AI tends to pay off first

Across a lot of small businesses, the same handful of workflow types keep turning out to be the best first moves. None of them are exotic. If you recognize your own week in here, you're probably looking at a starting point.

Moving information from one place to another. This is the quiet tax on almost every business - someone reading a value off an email, a PDF, or a form and typing it into another system. Invoices re-keyed into accounting. Order details copied from one tool to the next. Intake forms turned into CRM entries by hand. It's slow, it's mind-numbing, and it's exactly the kind of same-shaped, high-frequency work AI handles well, with a person spot-checking until the trust is earned. When people say AI gave them their afternoons back, this is usually what they mean.

Drafting the replies you send over and over. Support questions, quote requests, appointment confirmations - the same five customer emails you could nearly type in your sleep. AI is good at the first draft. It gets most of the way there in seconds, and a person edits and sends. You're not removing the human. You're removing the blank page. For a team that lives in its inbox, that alone can be the difference between replies going out before lunch and after.

Summarizing long things into short ones. Contracts, reports, application packets, long email threads, a stack of meeting notes. Reading to find the three things that matter is slow human work, and it's the sort of thing AI does quickly - give me the key terms, flag anything unusual, tell me what changed since the last version. You still read the full document when the stakes are high. But for triage - what is this, and does it need my attention? - a good summary saves real time.

First-pass research. The background you pull together before a sales call. The quick comparison of three options. The "what do we actually know about this company" digging. AI won't replace your judgment, and you should never trust it blindly on facts. But it's a strong research assistant for the first pass - gathering, organizing, and laying out what's known so you start from a draft instead of a cold screen, then check what matters.

Notice the thread running through all four. The AI takes the first, fast, dull pass; a person takes the last, careful one - the judgment, the fact-check, the send. That division of labor is why these work. They don't ask the technology to be perfect or to use judgment it doesn't have. They ask it to take the tedious first pass off a person's plate, and that's a job it can actually do.

Where it usually isn't worth it yet

Honesty is the whole point of a map - it has to show the dead ends too. Here are four kinds of work where, for a small business right now, AI tends to disappoint or cost more than it saves. None of these are forever. They're "not yet."

Work where being wrong is expensive and hard to catch. Anything that goes straight to a customer or a regulator with no human in between - legal answers, medical or financial advice, final numbers on a tax filing. AI sometimes produces a confident, well-written answer that happens to be wrong, and if no one is positioned to notice before it lands, the occasional miss can cost far more than the task ever saved. The same caution applies to work that has to be exactly right every single time. If 99 right out of 100 isn't good enough, this isn't the place to start.

Work that runs on judgment, relationships, and trust. The hard conversation with an unhappy client. Whether to hire someone. How to handle a delicate negotiation. These lean on context, read-the-room instinct, and accountability that a tool simply doesn't have. AI can help you prepare - draft talking points, lay out the options - but the call itself is yours. Handing it over doesn't save time; it costs you the thing that makes you worth dealing with.

Rare and one-off tasks. The setup cost of any AI workflow - describing it, connecting the tools, checking it until you trust it - only pays back when the task repeats. Automating something you do three times a year almost never earns back the effort. Do those by hand. Spend your setup energy on the things that happen daily.

Work built on messy or scattered data. If the information a task depends on lives in someone's head, is spread across six spreadsheets that don't agree, or is riddled with errors, AI won't fix that - it will faithfully repeat the mess, faster. The old rule holds: garbage in, garbage out. Sometimes the honest first project isn't an AI project at all; it's tidying the data so a later one can work. That's not a detour. That's the foundation.

If your best candidate sits in one of these four, it's not a failure of imagination on your part. It's a sign to either pick a different starting point or do the boring groundwork first. Both are smart. Neither is what a vendor will tell you.

Buy, build, or wait

Say you've found your task. There are really only three ways to get it done, and knowing which one you're in keeps you from overspending or overbuilding.

Buy when a tool already does the thing well and fits the software you already use. For common needs - transcription, basic document reading, drafting help inside an app you already own - someone has built a good, supported product, and buying it is faster and cheaper than anything custom. This is the right first move more often than people expect. Try the off-the-shelf option before you build anything.

Build when the task is specific to how you work and no off-the-shelf tool fits the way your business actually runs. Building, for a small business, rarely means inventing AI from scratch - it means connecting the tools you already use and wrapping a bit of AI around your own process. More setup, more payoff when it fits, and worth it when the task is frequent and genuinely yours. This is the part we spend most of our time on, and it's where a careful outside hand usually pays for itself.

Wait when the data isn't ready, the stakes are too high to start safely, the task is too rare to bother, or the tools just aren't good enough yet for your case. Waiting is a real, respectable decision - not a failure of nerve. The cost of waiting six months on the wrong project is almost always lower than the cost of rushing it. Put it on a list, set a date to revisit, and spend your energy where it pays off now.

Most small businesses we work with end up with a mix: buy the common stuff, build the one or two things that are truly theirs, and consciously wait on the rest. The skill isn't picking one philosophy. It's matching each task to the right answer.

How to know in two weeks if you picked right

The last piece of the map is how you'll know, quickly, whether the move was right - before you've sunk real money into finding out. You don't need a quarter. You need about two weeks and one number you already track.

Start by writing down the before. How long does the task take today, how often does it happen, and who does it? Don't trust memory - memory flatters. If you have to, time it honestly for a few days first. Then pick one number that should move if this works: hours per week on the task, replies sent before lunch, days to get a quote out, errors caught before they reach a customer. One number, not five.

Then run the new way next to the old way for two weeks. Keep doing the work the way you always have, and run the AI version alongside it. That protects you if it flops, and it gives you a clean comparison - same task, same weeks, old way versus new, measured against that one number. Keep a person checking the output the whole time. You're testing whether the assisted version is faster and at least as good, which is a lower and far more useful bar than "can we remove people."

Here's the part almost everyone skips, and it's what makes the whole thing honest: decide before you start what would make you stop. Write down both directions. Success might be "the number moved by a third and quality held." Stop might be "the number didn't budge," or "quality slipped," or "it took more

babysitting than it saved." If you hit the stop condition, you stop.

And this matters - stopping is not failure. A pilot you're allowed to stop is a pilot you can run honestly. If two weeks in the number hasn't moved, you've spent very little and learned something real about where AI fits your business and where it doesn't. That's the cheapest research you'll ever buy. Useful either way.

A clear next move, useful either way

If you'd like a second set of eyes before you commit, that's exactly what our first call is for. It's thirty minutes, it's free, and there's no slideshow. Tell us how your week actually goes, and we'll help you spot the one or two tasks where AI would earn its keep first - or tell you honestly if the right answer for now is "wait." Either way you'll leave with a clear next move you can act on, with us or entirely on your own.

No pressure, no jargon, no pitch for things you don't need. We make our living building the handful of automations that pay for themselves and handing them off so your team runs them without us. The fastest way to find out whether one of those is hiding in your business is a short conversation.

Call us at **512.234.5665**, or work through the worksheet on the next page first and bring it with you. Most people find that the act of filling it in answers half their questions before we ever talk. The map is yours to keep - weeks, not months, and useful either way.

Appendix: Your AI Starting Map worksheet

Print this page. Fill it in by hand, ideally with the one or two people who know the work best. Twenty minutes here is worth more than twenty articles about AI.

Step 1 - List your candidates

Write down 5 to 8 tasks that eat time in your week. Don't overthink it - the ones that made you sigh this week are perfect.

- 1 _____
- 2 _____
- 3 _____
- 4 _____
- 5 _____
- 6 _____
- 7 _____
- 8 _____

Step 2 - Score each one

Score every task 1 to 5 on each question. Higher is always better - a 5 means "great first-project quality." Add the row across for a total out of 20.

Scoring the last column: **5** = a mistake is cheap and easy to catch. **1** = a mistake is costly or hard to catch.

Task	Happens often? (1-5)	Same shape every time? (1-5)	Disliked? (1-5)	Safe to be wrong? (1-5)	Total /20

Step 3 - Find your top one or two

Circle the one or two highest totals. Those are your starting candidates. A score of 16 or higher, with no single column down at a 1 or 2, usually means "start here." If your top task scored low on "safe to be wrong," look at your next one instead - or plan to add an extra human check before you begin.

My top candidate: _____

Runner-up (the likely second project): _____

Step 4 - Write the first project

Keep it to one task. If you can't describe it in a single sentence, it's too big.

- The one task we'll start with: _____
- Who runs it (one person): _____
- The single number we'll watch: _____
- That number today (measured, not guessed): _____
- What "working" looks like in two weeks: _____
- What "not working" looks like - our stop signal: _____
- Start date: _____ Two-week check-in date: _____

Want a second opinion on what you wrote? Bring this page to a free call: **512.234.5665**. Thirty minutes, a clear next step, useful either way.