The SMB AI Agent Playbook

Austin Moore, AI Strategist

Moore Agentic LLC | Tampa Bay, FL

mooreagentic.com

A Paradigm Shift in Software: What Every Business Owner Needs to Know Before Building Their First AI Agent

Table of Contents

A Note Before You Start

I wrote this because I kept having the same conversation with business owners.

Someone would come to me after spending $30,000 on an AI project that never worked — or worse, they'd have a system running that nobody actually used. And when we'd trace back where things went wrong, it was almost never a technology problem. It was a strategy problem. They built before they understood. They bought what sounded impressive instead of what they actually needed.

I'm an AI strategist, which means my job is to help businesses figure out where AI actually makes sense — and just as importantly, where it doesn't. I'd rather have an honest 30-minute conversation than watch you spend six months building something that doesn't move the needle.

That's the spirit of this guide. **Read it before you spend a dollar.**

01

What AI Agents Actually Are (And What They're Not)

Let's start with the definition that matters, not the marketing version. An AI agent is software that can observe a situation, reason through what to do next, take an action, and then adapt based on what happens. It's not just a chatbot that answers questions. It's a system that can pursue a goal across multiple steps — without a human approving each one.

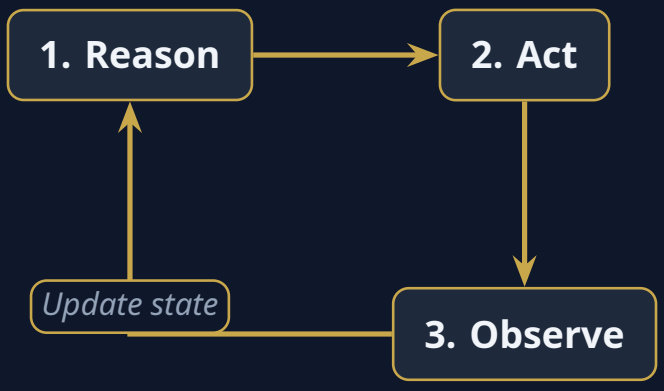

Here's a **concrete example**. A standard chatbot can answer "What's your return policy?" An AI agent can take a customer's return request, look up their order in the CRM, verify whether it falls within the return window, process the refund, send a confirmation email, and log the outcome — all in one go, without a human touching it. That's the difference. **A chatbot responds. An agent acts.**

> **What AI agents are not:**
>
> - **They are not magic.** An AI agent is only as good as the data it has access to, the instructions it's been given, and the tools it can use. A poorly configured

agent will make bad decisions confidently.

- **They are not a replacement for clear processes.** If your team doesn't have a consistent, documented way of doing something today, an AI agent won't fix that. It will automate your chaos.
- **They are not autonomous employees.** Even the most capable AI agent needs guardrails, monitoring, and occasional human review — especially when it's new. Think of a first deployment like a new hire's first week.

Why this matters for your business:

The businesses that get the most from AI agents are the ones who go in with clear expectations. They know what they're automating, they know what success looks like, and they've thought about where a mistake would hurt. That clarity is what separates a project that delivers ROI from one that becomes a cautionary tale.

02

The 3 Types of AI Agents Every SMB Should Know

Not all AI agents do the same thing. Choosing the wrong type for your use case is one of the most common mistakes I see. Here are the three categories that matter most for small and mid-sized businesses.

Type 1: Conversational Agents

These are agents that interact with people — customers, prospects, or your own team — through natural language. They read incoming messages, figure out what the person needs, and respond or take action.

Use cases: customer support, lead qualification, internal help desk, FAQ handling, appointment scheduling.

Strategic Note: These live or die by the quality of their knowledge base. If you feed them incomplete or inaccurate information about your products, policies, or processes, they'll give bad answers — and they'll do it politely, which makes it worse.

Type 2: Task Agents

These agents work in the background, executing specific, defined workflows without being prompted by a user in real time. You set them up to watch for a trigger — a new form submission, an incoming email, a record update — and they take it from there.

Use cases: lead enrichment, invoice processing, report generation, data entry and cleanup, monitoring and alerts.

Strategic Note: Task agents are where most SMBs find their fastest ROI. The workflows they automate tend to be repetitive, rule-based, and time-consuming — exactly the kind of work that frustrates good employees and is expensive to do manually at scale.

Type 3: Orchestration Agents

These coordinate other agents. They receive a complex goal, break it into sub-tasks, and assign each sub-task to a specialized agent. Think of it like a project manager delegating to a team.

Use cases: complex sales workflows, multi-step research and reporting, coordinating across departments, building products that require multiple AI capabilities working together.

Strategic Note: Orchestration agents are the most powerful — and the most complex to build and maintain. For most businesses just starting with AI, this is not where you begin. It's where you grow into.

Which type do you need?

Start by asking: is this about talking to someone, doing a task, or coordinating multiple things? Most SMBs need one or two well-built task agents before they need anything more sophisticated. Get the basics working first.

03

How to Pick Your First Use Case (Without Wasting Money)

The biggest mistake in first AI projects isn't a technology mistake. It's choosing the wrong problem to solve. I've seen businesses transform their operations with a simple agent that saved each employee 90 minutes a day.

The Four Questions:

1. **Is this process repetitive and rule-based?** AI agents are exceptional at tasks that follow a consistent pattern. If the work requires nuanced judgment that changes significantly case by case, proceed with caution.
2. **Does it happen often enough to matter?** Automating a process that happens twice a month won't change your business. Automating something that happens 50 times a day will.
3. **Is there a clear, measurable outcome?** The best first AI projects have a result you can count: hours saved, cost per transaction, time to respond. Avoid use cases where success is "fuzzy."
4. **What happens when it makes a mistake?** Every AI agent will make mistakes. Choose something where errors are low-stakes — where a human can easily catch and fix them before they cause damage.

Red Flags:

If your first answer to "what should we automate?" is something customer-facing and high-stakes — like handling complex client complaints or making credit decisions — **slow down**. Build confidence with an internal process first. Prove the value, learn how the system behaves, then expand to higher-stakes applications.

A Practical Starting Point:

List the five tasks in your business that are most repetitive, most time-consuming, and most rule-based. Score them against the four questions. The one that scores highest is your first project. Don't overthink it. A good first project is better than a perfect one.

04

Build vs. Buy: What Makes Sense Under 100 People

You've identified a use case. Now you face the question: do you build a custom AI agent, buy an off-the-shelf tool, or hire someone to do it for you?

When to buy (off-the-shelf tools):

If your use case is common — scheduling, basic customer support, document summarization — there's a good chance a product already does it well. Tools like Zapier, Make, Intercom, and dozens of others have AI built in now. Buying is almost always faster and cheaper than building.

The test: Can you describe your use case in one sentence? If yes, there's probably a tool for it.

When to build (custom development):

Custom agents make sense when your use case is specific to your business — when it requires your internal data, your specific processes, or your unique logic.

The Real Estate Example: The real estate reporting tool I built for a client saved 12 hours a week because it pulled from their internal CRM, matched data against local market comps, and formatted reports exactly the way their agents needed them. No product on the market did all of that. Building made sense because the value was in the specificity.

The test: Can you describe exactly what the agent needs to know and do that's unique to your business?

When to bring in an AI strategist:

If you need something custom but don't have the technical team, a strategist is the right call. You are paying for the expertise to design the system correctly and avoid the expensive mistakes. They should ask more questions than they answer in the first meeting and have real examples of work they've done — not just capabilities they claim to have.

> **The Hidden Cost of Building:**
>
> Building a custom AI agent isn't a one-time expense. There's ongoing maintenance, prompt refinement as models update, monitoring, and rebuilding as your needs change. Budget for the full lifecycle, not just the initial build.

05

How to Measure AI ROI Before You Spend a Dollar

One of the most common mistakes I see is businesses approving AI projects without a clear definition of success. You can avoid this entirely by doing the math before you build.

The ROI Framework — Three Types of Value:

- **Time savings** — Total manual hours per week × fully-loaded cost per hour. If an agent handles 80% of that work, you've got your baseline savings estimate.
- **Revenue impact** — Does this process affect sales? Faster lead response time or higher-quality proposals affect close rates. Even conservative conversion improvements often lead to compelling numbers.
- **Error reduction** — What's the current error rate and cost (rework, refunds, churn)? An agent that eliminates errors in high-volume processes pays for itself in the first quarter.

The Math:

Estimate value across these three categories vs. the cost to build and maintain. If the payback period is under 12 months, it's almost certainly worth doing. **Most well-scoped SMB AI projects pay back in 3–6 months.**

What to measure once you're live:

Pick two or three metrics: hours saved, response time, error rate. Establish your baseline now. Then measure the same thing after 30, 60, and 90 days. If the numbers aren't moving, the agent needs tuning, or the use case wasn't as valuable as you thought.

A Note on Intangibles:

There are real benefits that are harder to quantify — employee morale, faster customer experiences, and data quality improvements. These are real, but don't let them become a substitute for hard numbers. If you can't show the math, you can't justify the investment.

06

Your 90-Day Roadmap to Your First Production Agent

Most AI projects fail not because the technology doesn't work, but because the rollout doesn't.

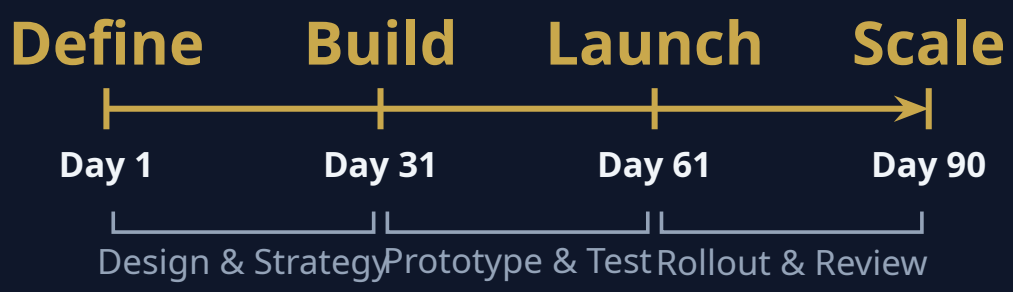

Days 1–30: Define and Design

Week 1–2: Document the process in detail. Standardize it first.

Week 2–3: Define success metrics and establish your baseline.

Week 3–4: Design agent architecture — tools, data, and human review points.

No code yet. The time you spend here saves weeks of rework later.

Days 31–60: Build and Test

Week 5–6: Build with small, controlled scope (80% routine cases).

Week 7–8: Test with real data (not live customers). Find edge cases. Tune prompts, logic, and guardrails. This phase should feel slow.

Days 61–90: Launch and Learn

Week 9–10: Soft launch to 20% of volume. Monitor agent decisions daily.

> Week 11–12: Expand to full volume. Debrief with the team and collect feedback.

The Most Important Thing:

Pick one person to own this. Not a committee, not IT as an afterthought. Someone who cares about the outcome, understands the business process, and has authority. AI projects with a clear owner succeed at much higher rates.

Technical Appendix: The Infrastructure of Intelligence

Derived from the Google Cloud Startup framework, these components ensure your agent is production-ready.

ReAct Architecture:
The ReAct loop (Reason + Action) assessment allows for synergistic intelligence where reasoning helps the model update action plans, and actions gather information to inform reasoning.

Grounding via Agentic RAG:
Retrieval-Augmented Generation transforms an agent from a passive model into an active problem solver. It ensures the agent retrieves "Ground Truth" data before generating an answer.

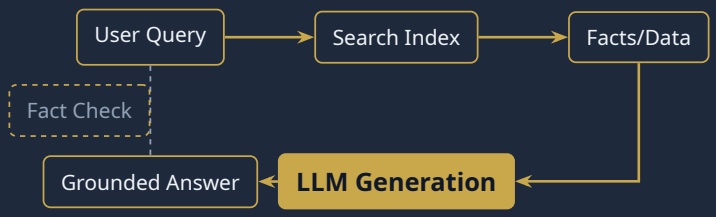

What Comes Next

If you've read this far, you know more than most business owners who are evaluating AI. You understand what agents actually do, how to choose a use case, and how to evaluate your options.

The next step is a conversation.

I offer a free initial call to any business that's seriously evaluating AI. No pitch, no pressure. We'll talk through where you are and whether an AI agent is actually the right tool for it.

If you're ready, you know where to find me.

Austin Moore
Moore Agentic LLC | Tampa Bay, FL
mooreagentic.com

This playbook was written based on real experience building AI systems for businesses. Technical frameworks are based on Google Cloud's Agent Development Kit (ADK) and AgentOps standards.