

One Person, Funded-Team Output

The honest ledger of what one operator directing AI actually shipped in one month — including the rejections, the wrong claims that reached production, and the backup that turned out to be dead. The leverage is real. So are the failure modes.

A build-in-public accounting, not a hype piece

THE TRUST BOUNDARY + THE SOLO-LEVERAGE WORKSHEET INCLUDED

BY ADAM FISCHER · PRAXIS-HQ.AI · THEORY INTO PRACTICE

What's inside

- 1 The month, in one table
- 2 What "directing AI" actually means
- 3 The real cost accounting
- 4 Where it broke — the failures, unedited
- 5 The trust boundary — what never gets delegated
- 6 The replication playbook
- 7 Where this model breaks down
- 8 The worksheet — score your own solo leverage

Every claim in this paper is checkable against public records — App Store listings, a live website, published papers, git history. Where a number is private, it's stated as private rather than invented. That discipline is not a formality; it's the subject of section 4.

1 The month, in one table

Praxis is a one-person company with no employees and no contractors. I direct AI systems that do the building; I make every decision and answer for every release. Here is what that arrangement actually shipped in July 2026 — not projections, shipped, checkable:

Shipped in July 2026	The particulars
An iPhone, iPad and Apple Watch product, live	Roofing Expert — subscription roof-inspection reports on iPhone, iPad and Apple Watch. Version 1.0 cleared App Review July 7; by July 22 it was on version 1.6, with 1.7 submitted the same week.
Two more apps, spec to submission in one week	LossList (homeowner claim documentation) and ClaimForge (public-adjuster workbench, iOS + macOS): specification locked July 15, the full three-app data loop running as code by July 21, field-tested via TestFlight and submitted to App Review July 22.
An eleven-paper technical library	Eight volumes written at the start of the month, a flagship and three more by month's end — published on praxis-hq.ai and Gumroad. The most recent went topic-to-published in roughly a day.
A company site, live	Multi-page site at praxis-hq.ai — products, consulting, publishing library, team — with a separate audience-specific landing page for the trades.
A consulting practice, built to gate	The Digital Exposure Audit: service spec, delivery playbook, authorization agreement, and pricing — deliberately held before the first paying client until insurance and contracts are in place.
Outreach infrastructure	A researched 68-prospect campaign list (every address verified against the company's own site, every domain checked), email templates, and a send playbook — built, and gated behind the apps going live.
The operating stack	Brand, custom-domain email, e-commerce, legal pages, App Store Connect operations for three products, versioned backups on multiple legs.

What would that list cost from a funded team? I'm not going to pretend to know your market's salaries, and this paper doesn't invent numbers — run your own: price a senior iOS developer,

a web developer, a technical writer, and an operations hire for a month against the table above. The comparison isn't subtle either way.

The claim is not "AI did this." The claim is: one experienced operator, directing AI, produced the output of a team — and made every mistake a team's process exists to catch. Both halves are the story.

2 What "directing AI" actually means

The phrase sounds like prompting tricks. It isn't. After a month of running a company this way, the actual job description is older than the technology — it's operations management, compressed:

- **Specification.** The AI builds fast and builds wrong if the destination is vague. The operator's leverage is knowing the domain well enough to specify — in my case, knowing what a carrier-ready roof report must contain because I've stood on the roofs and sat at the adjuster's desk. Domain expertise didn't get less valuable; it became the bottleneck asset.
- **Judgment.** Dozens of times a day: is this right, is this good enough to ship, is this the correct trade-off? The AI proposes; the operator disposes. Nothing in the table above shipped without a human decision.
- **Verification.** The discipline of not believing your own tools — insisting claims be checked against reality before they're published. Section 4 is a tour of what happens at the exact moments this discipline slips.
- **Saying no.** The highest-value interventions of the month were refusals: features held because a statute made them risky, copy killed because it overclaimed, a legal question routed to a lawyer instead of answered confidently by a machine.
- **Written memory.** The whole arrangement runs on files — state, decisions, rules, lessons — that let any working session pick up exactly where the last one stopped. (That method is Volume 10 of this series; it is not optional at this pace.)

Notice what's absent: coding skill. I don't have a programming background — thirty years of technology *operations*, yes, but the leverage here isn't writing code, it's running a shop where something else writes the code and everything still has to meet spec.

3 The real cost accounting

I am going to disappoint anyone who came for a spreadsheet. There is a version of this section that itemizes every subscription to the dollar, and I have decided not to write it — partly because a competitor's cost sheet is the least useful thing you can copy, and mostly because the itemized version invites you to draw exactly the wrong conclusion.

The shape is what matters, and the shape is this: **one professional-tier AI plan dominates the entire fixed stack**. Everything else — the developer program, a domain, static hosting, a storefront that takes its cut per sale rather than a fee — is rounding error beside it. There is no payroll, no office and no benefits line, because there is no one to pay.

Two honest footnotes to that, and they matter more than any number I could print. First, **the real cost is time** — this month was not a part-time effort, and no tooling makes judgment-hours free; what changed is what an hour of judgment now produces. Second, **near-zero fixed cost is a launch posture, not a destination**: the consulting line's compliance gate (insurance among it) and the growth items ahead all add real costs, on purpose, when revenue justifies them. The discipline isn't spending nothing — it's adding fixed costs only behind evidence.

So the useful takeaway is not a budget to copy. It is that the tooling was never the constraint, and if you are waiting to be able to afford this way of working, you are already able to.

4 Where it broke — the failures, unedited

This section is the reason to trust the rest of the paper. A month at this pace produced real failures, and they cluster into a pattern worth more than any success story: **AI fails confidently, and the failures survive until a verification habit catches them**.

Wrong claims reached production copy

Marketing text went live describing my background in a way that inflated it — a plausible-sounding line, repeated by the AI across pages and assets because early drafts contained it. Worse: a "we checked everywhere" sweep declared it cleared while copies survived on live pages the sweep never actually reached, because the search tooling silently skipped a directory. It took a deliberate re-audit, with the blind spot fixed, to kill every instance — including two

discovered days after the first "all clear." The lessons, now house rules: exact claims only, verify the *generator* and not just the output, and never trust a sweep you haven't proven reached everything.

Two App Store rejections in one evening

Both companion apps were rejected by the same automated check — a required legal link missing from store metadata. Cost: an evening. It was fixed and resubmitted within hours, and the rule went into the submission checklist so it can't recur. Cheap as failures go, but it makes the point: process gaps at a one-person company have no second pair of eyes to catch them — the checklist is the second pair of eyes.

The runaway render

A document-generation approach that had worked elsewhere hit an edge case and silently wrote gigabytes of garbage before being killed. The replacement pipeline got a watchdog that aborts loudly instead. Lesson: at this speed you will occasionally automate a mistake at scale; the answer is tripwires that make failure loud, not hope that it stays small.

The backups that weren't

An audit of the backup system found one leg (a machine-level backup) dead since a move months earlier, and another quietly accumulating duplicates because its cleanup step bailed out exactly when it was needed. Everything looked fine until someone actually checked. Both fixed the same night — and the general lesson is the month's theme in miniature: **an unverified safety system is a belief, not a system.**

THE PATTERN ACROSS ALL FOUR

None of these were exotic. Every one was ordinary operational entropy — moving fast, trusting a tool, assuming a check had run. The difference between this model working and not working isn't avoiding such failures; it's the standing habits that surface them in hours instead of months: sweeps that prove their own coverage, checklists that encode every burn, tripwires on automation, and audits of the boring infrastructure nobody wants to look at.

5 The trust boundary — what never gets delegated

The month settled a question I now consider the core design decision of a one-person AI company: *where does delegation stop?* The boundary that emerged, drawn in places where a mistake is either irreversible or lands on a human being:

- **Public claims about people.** Anything stating my background, credentials, or history is verified by me, word by word, against reality — after watching plausible inflations propagate. Titles that imply licensure get special scrutiny, because in regulated industries a title is a legal claim.
- **Legal and sworn anything.** Contracts, legal pages, statutory questions, anything touching a court or a regulator: AI drafts and researches, licensed humans decide. Several of the month's best decisions were the AI itself flagging "this one goes to the attorney."
- **Money and commitments.** No payment, purchase, submission, or send happens without my explicit go. Drafts are automated; the last click is not.
- **The regulated line in my own industry.** The claims apps are deliberately engineered so the software never interprets policies or promises outcomes — features have been kept out of the roadmap, or moved between products, specifically to stay on the right side of licensure law. The AI proposes features; the boundary disposes of some of them.

Delegate the production. Never delegate the accountability. Every failure in section 4 got caught precisely at this boundary — and everything inside it moves at machine speed because the boundary holds.

6 The replication playbook

What would you actually need to run this model? Four things, and only one of them is software:

1. **Deep domain knowledge in something.** The model amplifies expertise; it doesn't substitute for it. My leverage came from knowing exactly what a roof inspection must produce and what a claim file must survive. Yours is wherever you've spent your years. If you can specify what "right" looks like in your field, you have the scarce input.

2. **A professional-tier AI setup you'll actually live in.** Not the free tier — the versions that hold context, use tools, and work in long sessions. One primary system, learned deeply, beats five dabbled with.
3. **The written-memory discipline.** Context files, a decision log, a state file, a lessons file (Volume 10). This is what turns disconnected sessions into a compounding operation. It costs minutes a day and it is the difference between a tool and a company.
4. **An operating cadence with verification built in.** A rhythm — what happens daily, weekly, per release — where checking is a step, not a virtue. Section 4 is what happens in the gaps between checks; the cadence is how you shrink the gaps.

Then the sequencing advice, learned rather than theorized: **ship one real thing all the way through before you diversify.** Every capability the later products reused — store operations, document pipelines, legal pages, release discipline — was earned end-to-end on the first product. The second and third moved fast *because* the first moved carefully.

7 Where this model breaks down

An honest paper names its own limits. From the inside, four are visible:

- **The bus factor is one.** Every escalation path ends at the same person; there is no colleague to catch what I miss when I'm tired, and section 4 shows what slips through. Mitigations exist — written memory, checklists, the advisory bench being assembled — but a mitigation is not a colleague.
- **Judgment doesn't parallelize.** The AI can draft five things at once; I can only *decide* one at a time. At scale, decision quality — not production — is the ceiling, and the honest end-state for a growing solo-AI company is hiring for judgment, which returns you to the org-chart questions this model deferred.
- **Some work is legally human.** Licensure, representation, signatures, insurance — whole categories where the constraint isn't capability but law, and correctly so. A one-person AI company in a regulated industry spends real energy staying on the right side of lines that a bigger firm staffs for.
- **Marketing and distribution still run on human time.** Building the campaign took days; the relationships it's aimed at will take months. AI compresses production spectacularly and persuasion only modestly. The table in section 1 is a supply-side story — demand is earned at human speed.

8 The worksheet — score your own solo leverage

Score each row 0 (not true of me) to 2 (solidly true). It's a mirror, not a gate.

Statement	0–2
There's a domain where I can specify what "right" looks like better than almost anyone I know.	
I can describe work precisely in writing — spec, acceptance, edge cases.	
I'm willing to verify outputs against reality even when they look right.	
I keep written records of decisions and lessons without being forced to.	
I can tell the difference between confident and correct.	
I know which parts of my field are regulated, and where the licensed lines are.	
I can sustain an operating rhythm alone — daily, weekly, per release.	
I have, or will build, the financial runway for the months before revenue.	

Reading it: the first two rows are the engine — without specification ability, the model produces fast wrongness. Rows three through five are the safety system; section 4 is what their absence looks like. The last three are the company around the engine. Low scores aren't disqualifying; they're your build list, and every one of them is buildable.

The close: the technology finally caught up to a very old idea — that one person who deeply understands a problem, given sufficient leverage, can build the solution without permission, funding, or a team. The leverage is now real. What it amplifies is still you: your expertise, your judgment, and — as a month of receipts shows — your mistakes. Bring the discipline, and the output of a funded team is genuinely on the table.

Direct the machine. Own the outcome. That's the whole paper.