



THE CASE FOR EMBEDDING

Buy the expertise, not the headcount.

Why the fastest way to ship AI software is to embed a studio that has already made the mistakes you're about to make.



THE CASE

AI made building cheap. It made shipping harder.

Anyone can produce a working prototype now. The scarce skill moved somewhere else: to the judgment about what's worth building, and to the last-mile engineering that turns a demo into something a customer can depend on.

That expertise is hard to hire. The field is two years old, the people who have shipped real AI to production are rare, and by the time you've run a search and closed a senior engineer, the tooling has moved again.

So teams stall. Not for lack of ideas or effort, but because the specific experience of getting AI software over the line lives in very few heads. Embedding a studio borrows those heads, on your codebase, now.

3–6 mo

to run a search and close one senior AI engineer, before they've written a line for you.

ILLUSTRATIVE

weeks

for an embedded team to be productive on your code, with the judgment already in the room.

ILLUSTRATIVE



WHAT YOU GET

Expertise you can't hire fast enough.

Eight capabilities that usually take a team years to assemble, working on your product from week one.

Agentic engineering

Claude Code, agent harnesses and agentic workflows used to build, not just to demo.

Evals & observability

Know whether the model is right, and see it when it isn't, in production.

System design

Architecture that holds up when the prototype meets real load and real data.

Retrieval & RAG

Grounding models in your data without the accuracy falling apart at scale.

Prompt & skills

Prompt engineering and reusable skills that make behaviour predictable.

CI/CD & testing

Automated testing and delivery so shipping is boring and reversible.

Production hardening

Auth, permissions, secrets, rate limits and cost caps done once, done right.

AI product design

Interfaces that make a probabilistic system feel trustworthy to real users.



YOUR OPTIONS

Hire, outsource, or embed.

Each is right for a different situation. Here is the honest version of the trade-offs.

	Hire in-house	Outsource a project	Embed a studio
Time to productive	3–6 months to search and close	Fast start, slow to learn context	Weeks, on your codebase
Judgment on what to build	Grows over years	Usually out of scope	The core of the engagement
Knowledge afterwards	Stays, if they stay	Leaves with the vendor	Left with your team
Best when	Steady, long-term load	A well-specced one-off	You must move now and keep it

If your load is steady and predictable, hire. If the work is truly one-off, outsource it. Embedding earns its place when you need to move at once and own what's left behind.



HOW WE WORK

The people who scope it are the people who ship it.

No hand-offs, no re-learning tax, no gap between the deck and the code. We join where you already are and stay until it's launched.

01 Start wherever you are

A rough idea or something that already half-works. We meet the project at its real state, not a tidy version of it.

02 One continuous team

The same people scope, prototype, harden and launch. Context compounds instead of leaking at every hand-off.

03 Leave a team that can carry it

We hand over a product your people can run and extend, and the habits that keep it shippable after we've gone.

"We move in weeks, not quarters, because nobody has to re-learn the problem to make the next decision."



OUTCOMES

What changes when the expertise is in the room.

TYPICAL SHIFT, EMBEDDED VS. GOING IT ALONE

ILLUSTRATIVE

Time to first launch**faster****Projects that reach real users****higher****Cost per successful outcome****lower****Knowledge retained in your team****higher**

These bars show direction, not a promise. The point isn't a specific number; it's that every one of them moves the right way when the people doing the work have shipped this before. We instrument real engagements so the actual figures are yours to see.

BUILD WITH AI · LAUNCH WITH SUBWORKFLOW

Borrow the experience.

Embed a team that has already shipped AI to production, on your codebase, from week one.

See where we'd start.

Book a 30-minute discovery call. We'll tell you honestly whether embedding is the right call, and what we'd do first.

[Book a discovery call →](#)