

COMPARISON

Finn vs. building it in-house

The real cost of build vs. buy for production voice AI.

Why teams weigh it

Building a voice agent is a reasonable instinct — a working demo is a weekend project with a model API and a telephony provider. The question was never whether your team can build it. It's what it costs to run it in production, reliably, for years.

The tradeoff

The prototype is about 5% of the work. The other 95% is production: sub-second latency under real network jitter, telephony edge cases, patching when a provider deprecates an API, evals to catch regressions before customers do, and audited compliance. That's a standing team, not a sprint. Finn is that team — a maintained voice layer benchmarked on 50M+ minutes of production traffic — so your engineers stay on the product only you can build.

Side by side

Building in-house	Finn
Months of engineering to a working pilot	Live in under 30 days
You own latency, telephony, and uptime	Sub-400ms fused pipeline, 99.99% uptime, maintained
Re-work every time a model or API changes	Provider changes absorbed for you
Build SOC 2, HIPAA, GDPR, ISO 27001 yourself	Compliance built in and audited
Engineers off your roadmap	Engineers on your product

The build is a weekend. The maintenance is forever — and that is the part Finn owns.

Compare on your numbers — book a demo at www.hirefinn.ai/bookademo

Or start free at hirefinn.ai · hello@hirefinn.ai