

PRICING

Usage-based pricing (USD)

Pay per minute for what a call actually uses — model, voice, telephony. Free to start, no subscription, no minimums, per-second billing with no rounding.

How pricing works

- Priced by the component. Every piece — language model, text-to-speech, telephony — is priced openly, and you only pay for the pieces a call actually uses.
- Per-second billing, no rounding. You pay for talk time, not rounded-up minutes. No charge for SIP trunking or bringing your own telephony.
- Free to start, no lock-in. No subscription and no minimums on pay-as-you-go — monthly plans apply only beyond what your usage includes.

Two ways to buy

Pay-as-you-go	Enterprise
\$0.07–\$0.31 per minute, all-in	Custom pricing at volume
20 concurrent calls, more on demand	No cap on concurrent calls
Shared infrastructure	Dedicated stable server
SOC 2 Type II, HIPAA, GDPR, ISO 27001	+ BAA, custom MSA / DPA, SSO, RBAC
Community and email support	24/7 support + named account manager
Self-serve setup	Dedicated implementation support

Included on every plan

- Built-in fallback system, safety guardrails, and PII redaction — plus opt-out of recording and transcription per workflow or per call.

A typical 4–8 minute call runs well under \$1 at a ~\$0.12/min blended rate — a fraction of a fully-loaded human rep. See the full component rate card at www.hirefinn.ai/pricing.

Price it on your volume — book a demo at www.hirefinn.ai/bookademo

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