Recapp

Competitive Landscape

Who else already turns an in-person meeting into a transcript, a summary and a list of action items, what they charge for it, how long they make the ...

2026-09-01

COPUNDIT

Prepared by Copundit

AI for Advanced Ideation™

Who else already turns an in-person meeting into a transcript, a summary and a list of action items, what they charge for it, how long they make the user wait, and what the buyer of a 159 dollar pocket recorder was actually buying. Every company profile records the same three things a reader of this field asks for first: the form the product takes, how it captures the room, and the time from the end of a recording to the finished summary, or the fact that no such time is published.

Contents

Contents

Abbreviations

Abbr	Stands for	What it actually is (plain English)
AI	Artificial Intelligence	Software that produces text or decisions from data instead of following fixed rules
ARR	Annual Recurring Revenue	Subscription revenue for one year, counted as if the current month repeated twelve times
ASR	Automatic Speech Recognition	The software that turns recorded speech into written words
BIPA	Biometric Information Privacy Act	An Illinois law that treats a voiceprint as biometric data and sets damages per violation
CAC	Customer Acquisition Cost	What a company spends to win one paying customer
CIPA	California Invasion of Privacy Act	California’s recording law; it requires every party to a private conversation to consent
CNY	Chinese Yuan Renminbi	China’s currency; also written RMB
CRM	Customer Relationship Management	The system a sales team keeps its customer records in, such as Salesforce or HubSpot
ECPA	Electronic Communications Privacy Act	The United States federal wiretap law; it allows one party to a conversation to record it
GB	Gigabyte	A unit of storage; one gigabyte holds roughly a day of recorded speech at ordinary quality
GDPR	General Data Protection Regulation	The European Union’s data protection law; a voice recording of a person is personal data under it
HP	Hewlett-Packard	The personal computer maker that bought Humane’s assets and switched off its device
HQ	Headquarters	Where a company is based; recorded as “not established” where no source stated it
LLM	Large Language Model	A text-generating model; here, the part that writes the summary and pulls out the action items
MEMS	Micro-Electro-Mechanical Systems	The tiny chip-scale microphones used in phones and pocket recorders
MRR	Monthly Recurring Revenue	Subscription revenue collected in one month
OS	Operating System	The phone’s own software layer (iOS, Android) that decides which app may hold the microphone
RMB	Renminbi	China’s currency, priced here alongside its United States dollar conversion
SOC 2	Service Organization Control 2	An audited security report that corporate buyers ask a software vendor for before purchase
SSO	Single Sign-On	Corporate login: staff reach the tool with the company account instead of their own password
WER	Word Error Rate	Percentage of words a transcript gets wrong; the standard accuracy number for speech recognition

Field Map

This is consumer and small-team voice productivity software: products that listen to people talking and hand back written artifacts. The unit of output is the same across every company below and it is the same three things this project proposes to deliver, a **transcript**, a **summary** and a list of **action items**. There is no insurer, no employer benefit and no reimbursement layer anywhere in this field. The user subscribes through an app store's in-app purchase system, which takes **15 to 30 percent** of the subscription before the vendor sees it, or direct by card, so every price quoted below is a gross price and not revenue. Those paid prices have converged tightly, at **9 to 19 US dollars a month** for an individual tier. The free tiers have not: the 300-minute monthly allowance belongs to the hardware vendor and the bundled suites, and the five phone applications profiled here run from 30 minutes a month to unmetered capture with the summary or the history rationed instead.

The field divides by **how the audio gets captured and who owns the thing doing the capturing**, and that single axis explains most of what happens in it. First, the **dedicated recorder**, a small object with its own microphone, battery and storage, sold up front with a subscription behind it. Second, the **application on a phone the user already owns**, recording an in-person conversation through the phone's own microphones. Third, the **feature built into the phone by the company that made the phone**, free, with no download and no subscription. Fourth, the **feature bundled into software the buyer already pays for**, a suite or a mobile carrier, which reaches the user without ever appearing on a comparison page. Fifth, and included only as a neighbour, the **bot that joins a video call** and records the platform's own audio stream, which is a cleaner signal and a meeting nobody walks out of.

The central fact for anyone reading this with the project's thesis in mind is that **the second category is not empty and is not expensive**. Voicenotes does the exact interaction, one button on the device's own microphone, live transcript, summary and action items, at a free tier and 9 US dollars a month (https://voicenotes.com/pricing). Fathom gives unlimited recording and transcription away and caps only its summaries (https://www.fathom.ai/pricing). Google's Pixel Recorder and Apple's Notes ship the recording, the transcript and the summary pre-installed and free. And the hardware sold anyway, over **two million devices** worldwide by mid-2026 (https://kr-asia.com/tencents-rumored-plaud-deal-points-to-looming-ai-hardware-contest). The reason is not only inattention. In a convenience sample of 64 explicit purchase-rationale statements assembled from verified retailer reviews and hardware-owner forums, which by construction describes people for whom the hardware already won, the largest single category, **28.1 percent**, is capturing a cellular phone call, and neither platform lets any third-party application touch the microphone during a call (https://www.reddit.com/r/PLAUDAI/comments/1m5boj6/as_a_sales_rep_whos_tried_otter_and_notta_plaud/). Those buyers did not walk past an equivalent application; for that job there is none.

Two economic details shape the rest. The first is that **the device does not include the software**. At 1,200 transcription minutes a month, the two category leaders list the same annualised price, **8.33 US dollars per user per month**, Plaud Pro at 17.99 month-to-month and Otter Pro at 16.99 (https://www.plaud.ai/pages/plaud-ai-plan-pricing, https://otter.ai/pricing). Buy the annual plan on either one and the recurring bill is the same, because both annual plans land on that same 8.33 for the same 1,200 minutes, so removing the hardware removes a one-time premium and not the recurring bill. Setting the device's annual plan against the application's month-to-month plan is what makes the recorder look half price over a year, and that is a billing-cadence mismatch rather than a price difference. The second detail is that **the recurring bill is metered and the subscription is flat.** The assembled marginal cost of a recorded hour in the cloud is **0.21 to 0.62 US dollars**, dominated by transcription, with the summarisation call three orders of magnitude below it; at a 15 dollar subscription netting about 13, breakeven arrives at roughly **32.5 hours a month**, and a user recording two to three hours a working day costs the vendor more than they pay (https://www.assemblyai.com/blog/speech-to-text-api-pricing). That is what the two leaders' matching 300-minute free caps are for. One vendor inverts the whole model: **iFlytek** sells recorders at 140 to 299 US dollars with lifetime free on-device transcription and no subscription at all, which is an offer no application charging monthly can match.

A third number decides whether any of those prices works, and **nobody in this field publishes it**. Every traction figure in this document is an install count, a review count, a raise, a valuation or an annual recurring revenue figure, and not one of them is a usage figure: no weekly retention, no monthly churn, no share of installs that record a second meeting. That matters more here than in most subscription categories, because the meeting a person genuinely wants written up is rare, so the normal shape is a large install base sitting on a small habitual base, and it is the unstated reason every winner below takes its growth from team and enterprise seats rather than from individuals. The only figures available are borrowed from consumer productivity software generally rather than measured in this category: free-to-paid conversion of 2 to 5 percent, monthly churn of 5 to 8 percent, an acquisition cost of 20 to 40 US dollars and a lifetime value of 180 to 300. They carry no category evidence and are quoted here only because nothing better exists.

The market is also not the one the seed's demand evidence describes. **Plaud** is a United States entity headquartered in Delaware and San Francisco, founded by Chinese entrepreneurs on a Shenzhen supply chain, and its total 2024 domestic Chinese sales are estimated at **under 100,000 units against over two million shipped worldwide** (https://tracxn.com/d/companies/plaudai, https://kr-asia.com/can-plauds-overseas-first-strategy-give-it-an-edge-in-chinas-crowded-voice-recorder-market). The non-Western market is real, large and separate, characterised by roughly 42 US dollar Huaqiangbei white-label clones of the same card form factor and by free bundled suites. So the hardware demand this project reads as evidence is Western demand.

No product in this field is a regulated medical or safety device, so each profile below records **certifications and named litigation** where a regulator line would otherwise go. That surface is real and it is moving. Four United States class actions are pending against note-taker vendors and **none has resolved**: *In re Otter.AI Privacy Litigation* (Northern District of California, from *Brewer v. Otter.ai*, August 2025), *Cruz v. Fireflies.AI* (December 2025), *Basich v. Microsoft* (Western District of Washington, February 2026) and *Chamberlain v. Granola* (Northern District of California, July 2026) (https://www.workplaceprivacyreport.com/2025/08/articles/artificial-intelligence/ai-notetaking-tools-under-fire-lessons-from-the-otter-ai-class-action-complaint/, https://www.computerworld.com/article/4206255/granola-lawsuit-raises-concerns-over-ai-note-taking-app-privacy.html). The Granola case is the one that matters most commercially, because Granola processes audio on the user's own device and was sued anyway: the complaint attacks how the product markets covert capture, not where the compute runs. On-device processing is a go-to-market asset and it is not a legal shield.

Two numbers are worth knowing before reading the profiles. **Some vendors do publish a turnaround time and none publishes a guaranteed maximum.** Otter's help centre states 1.5 to 3 minutes to a finalized transcript and 2.5 to 6 minutes to a summary for a 30-minute meeting, hedged with peak-load language (https://help.otter.ai/hc/en-us/articles/360048322493-Transcription-processing-time-FAQ); Fireflies publishes 15 to 20 minutes (https://fireflies.ai/blog/how-to-use-fireflies-ai). **Nobody publishes any quality figure at all.** No vendor among the ones examined publishes a summary or action-item factuality metric, and the accuracy claims that do circulate (Notta at 98.86 percent, Speakwise at 95 percent or better) name no test set, no audio condition and no microphone, so they are marketing rather than measurement. Purchases in this market are being made on something other than measured output quality.

Active Players

Ordered by how directly each contests the buyer this project is going after: the person who would otherwise buy a dedicated recorder to capture a meeting they physically walk out of. The applications come first because they deliver the same three outputs from the same phone, so a customer choosing one of them is a customer not choosing this project. That five profiles below carry a direct-competitor tag is itself the finding: this is not a field where the product has to be imagined.

Applications that record the room through the phone's own microphones

These do what this project proposes, on hardware the user already owns, today, at prices from zero to 19 US dollars a month. The question they raise is not whether the application can be built but what an entrant's reason to exist is, given that all five ship it today, three carry venture funding, two are bootstrapped, and the free tier of one covers unlimited meetings gated only by how long the notes are kept. **Otter, Granola, Wave, Voicenotes, Fathom.**

Otter

Mountain View, United States | founded 2016 as AISense | about 70 million US dollars raised, last a 50 million Series B led by Spectrum Equity | DIRECT COMPETITOR

- **Website:** https://otter.ai/
- **Form:** app, on phone and desktop, plus a bot for video calls.
- **Capture:** both. For the in-person case, "mobile and desktop apps that record directly from the device microphone with no bot required", and the product claims to distinguish "voices in multi-person rooms" (https://otter.ai/blog/ai-notetaker-for-in-person-meetings).
- **Time to output:** published, and it is the only published figure in this document that names a recording length. For a 30-minute meeting recorded by Otter, the finalized transcript is typically available within 1.5 to 3 minutes and the summary within 2.5 to 6 minutes, with an explicit peak-load caveat and no guaranteed maximum (https://help.otter.ai/hc/en-us/articles/360048322493-Transcription-processing-time-FAQ). An independent 2026 review at unstated recording length measured 10 to 15 minutes in practice (https://hinoter.com/blog/10-best-ai-note-takers-in-2026-tested-compared).
- **What they do:** streaming speech recognition with speaker labelling, feeding a language model that writes a summary with decisions and assignable action items.
- **Product:** phone and desktop applications plus a web workspace, sold per seat.
- **Tech:** software only, running on the microphones of whatever device it is installed on, with cloud processing. The vendor claims "Live transcription with 95%+ accuracy"; that figure names no test set and no audio condition, so it is a marketing claim and not a measured word error rate.
- **Who pays:** the end user or the team. Free Basic at 300 transcription minutes a month with a 30-minute cap per conversation; Pro at 16.99 US dollars per user per month or 8.33 annualised for 1,200 minutes; Business at 20 to 30; Enterprise at 98 (https://otter.ai/pricing).
- **Who subsidizes:** none found. Venture funded.
- **Certifications and litigation:** SOC 2 Type 2, with a Business Associate Agreement available on enterprise tiers. Lead defendant in *In re Otter.AI Privacy Litigation* before Judge Eumi K. Lee in the Northern District of California, consolidated from *Brewer v. Otter.ai* (August 2025), pleading the California Invasion of Privacy Act and the Electronic Communications Privacy Act on the theory that the bot is an unauthorised third-party eavesdropper rather than a tool of the host; the privacy theory drew judicial skepticism toward the defence at an April 2026 motion-to-dismiss hearing and the matter is unresolved (https://uctoday.com/otter-ai-on-trial-and-the-ai-notetaker-industry-with-it). Model training on transcripts is opt-out by default.
- **Traction:** about 100 million US dollars of ARR as of March 2025 with 56.3 percent year-over-year growth (https://sacra.com/c/otter/).
- **Why they matter to us:** it is the largest company selling the seed's exact interaction, and its published turnaround is the number the walk-back promise is actually measured against by anyone who checks.
- **Their weakness:** growth comes from Business and Enterprise seats rather than from individuals, so the individual consumer it would have to defend is the part of its own business it monetises least; and the measured 10 to 15 minute wait sits outside the range its own help centre publishes.

Granola

London, United Kingdom | founded November 2022 | 192 million US dollars total, last a 125 million Series C at a 1.5 billion valuation | DIRECT COMPETITOR

- **Website:** https://www.granola.ai/
- **Form:** app, on macOS, Windows, iOS and Android.

- **Capture:** both. The company describes “capturing and transcribing in-person and virtual meetings” and the application “sits in the background on an iPhone or a MacBook” (https://finance.yahoo.com/news/ai-note-taking-app-granola-140000909.html). The iPhone application, which is the in-person path, launched April 2025.
- **Time to output:** the vendor’s own claim is enhanced notes within 1 to 3 minutes of the meeting ending, at an unstated recording length (https://www.granola.ai/blog/granola-pricing-privacy-tradeoff). No service commitment and no guaranteed maximum is published.
- **What they do:** listens in the background on the device, merges the user’s typed notes with a model-written summary, and keeps the resulting notes queryable as organisational context.
- **Product:** desktop and phone applications plus a workspace, sold per seat, with an enterprise tier added on the Series C.
- **Tech:** software only. Capture is local and bot-free; inference runs in the cloud through third-party models, and the audio is discarded once the text exists, which removes the storage liability and is why the company can offer unlimited free recording (https://www.granola.ai/security).
- **Who pays:** the end user or the team. Free with unlimited recordings but only 30 days of history; Business at 14 US dollars per user per month for unlimited history; Enterprise from 35 (https://www.granola.ai/blog/meeting-notes-tool-pricing-benchmarks). The free tier gates memory rather than minutes, one of the two free tiers in this document that do not meter capture at all; Fathom’s is the other, and it rations the summaries instead.
- **Who subsidizes:** none. Venture funded: a 4.25 million US dollar seed from Lightspeed and betaworks, a 20 million Series A led by Spark Capital in October 2024, a 43 million Series B led by NFDG at a 250 million valuation in May 2025 (https://sifted.eu/articles/note-taking-app-granola-raise), and a 125 million Series C led by Index Ventures with Kleiner Perkins at a 1.5 billion valuation, March 2026 (https://techcrunch.com/2026/03/25/granola-raises-125m-hits-1-5b-valuation-as-it-expands-from-meeting-notetaker-to-enterprise-ai-app/).
- **Certifications and litigation:** SOC 2 Type 2; the policy states third parties are not allowed to use customer personal data for model training, and in-room transparency features exist but are off by default. Defendant in *Chamberlain v. Granola* (Northern District of California, July 2026), a California Invasion of Privacy Act class action pleading wiretapping over silent device-level capture with no visible bot; unresolved (https://www.computerworld.com/article/4206255/granola-lawsuit-raises-concerns-over-ai-note-taking-app-privacy.html).
- **Traction:** revenue grew 2.5 times between January and March 2026, with named enterprise adoption at Vercel, Gusto, Asana and Mistral AI; absolute revenue is not disclosed (https://thenextweb.com/news/granola-series-c-meeting-ai-enterprise-context).
- **Why they matter to us:** two reasons and the second is the sharper one. The valuation is the number any claim that this category is unserved has to sit against; and the lawsuit against it sits on this project’s exact architecture, silent device-level capture with no bot in the room, which is the charter’s design and not an implementation detail.
- **Their weakness:** the centre of gravity is the knowledge-worker laptop and the enterprise context graph, not a pocket recorder. It asks the user to keep a note surface open, and it monetises memory rather than capture, which leaves the person who wants one meeting written up and nothing else on the free tier forever.

Wave AI Note Taker

United States (Mohrer Associates LLC) | founding date not established | bootstrapped, no venture funding | DIRECT COMPETITOR

- **Website:** https://wave.co/
- **Form:** app, on iOS, Android, Mac, Windows and the web.
- **Capture:** in-person via the device’s own microphones, mobile-first, with background recording that continues while the screen is locked.
- **Time to output:** not published.
- **What they do:** records a room through the phone, transcribes it, and returns a summary with action items as a separate artifact.

- **Product:** phone applications with desktop and web companions.
- **Tech:** software only, rented models, no hardware.
- **Who pays:** the end user or a small team. Free at 30 minutes a month; Pro at 11.67 US dollars a month billed annually for unlimited recording and transcription; Teams at 7.50 per user per month with a five-seat minimum (https://wave.co/pricing).
- **Who subsidizes:** none. Bootstrapped by an independent studio.
- **Certifications and litigation:** none named. No litigation found.
- **Traction:** 4.9 stars across over 12,000 iOS App Store reviews; a third-party tracker estimates 360,000 US dollars of MRR on 25,000 monthly downloads, which annualises to roughly 4.3 million (https://fluxbuilder.com/explore/wave-ai-note-taker-recorder/). That is one tracker's estimate and no audited figure exists.
- **Why they matter to us:** it is the existence proof for the bootstrap axis at exactly this interaction. A mobile-first application with no venture backing reached seven figures of annual revenue selling the same three outputs to the same person, which both validates the shape of the business and removes the "nobody has built it" premise.
- **Their weakness:** the tightest free tier in the category at 30 minutes a month, no enterprise motion, and a revenue figure that only a third-party tracker asserts.

Voicenotes

San Francisco, United States | founded 2023 | unfunded, bootstrapped | DIRECT COMPETITOR

- **Website:** https://voicenotes.com/
- **Form:** app, on iOS, Android and the web.
- **Capture:** in-person via the device's own system microphone. The product page states it "works for online as well as in-person meetings" (https://voicenotes.com/).
- **Time to output:** not published. The vendor's own line is "The moment the meeting ends, your summary is already waiting", which carries no number and no recording length.
- **What they do:** one button starts recording on the phone's system microphone, the transcript appears live, and an "Ask AI" function produces the summary and pulls out the action items on request.
- **Product:** a phone and web application; the recording, transcript and key points sit in one place.
- **Tech:** software only, no hardware, with bought-in speech and language models reported as OpenAI Whisper and GPT-4 Turbo (https://apps.apple.com/us/app/voicenotes-ai-notes-meetings/id6483293628). No vendor engineering page states the model stack, so treat the attribution as reported rather than confirmed.
- **Who pays:** the end user. Free at 100 weekly transcription minutes shared across meetings, memos and dictation with a 30-day history; Pro at 9 US dollars per user per month for unlimited minutes and history; Enterprise at 24, adding SSO and provisioning, consent management and no model training on customer data (https://voicenotes.com/pricing). The pricing page carries no lifetime tier.
- **Who subsidizes:** none found. The company is unfunded and bootstrapped (https://tracxn.com/d/companies/voicenotes).
- **Certifications and litigation:** n/a, not a regulated product. The enterprise tier advertises consent management and zero model training on customer data; no audited certification is named on the pricing page. No litigation against it was found.
- **Traction:** the homepage carries a "1 million downloads" badge, which is a cumulative install count and not an active-user figure; the iOS listing shows 4.8 stars from about 6,700 reviews. No independent user, download or revenue figure for this company could be found, and that absence is the single largest open question in this document: a product that does the whole job at nine dollars, whose reach nobody can measure.
- **Why they matter to us:** it is the clearest single answer to whether this product already exists. The exact interaction and the exact three outputs, on the phone, from a company with no funding, at a price a subscription cannot undercut.
- **Their weakness:** action items arrive only when the user asks for them through a chat query rather

than as a standing artifact after every meeting, and an unfunded team has no acquisition engine to put against competitors holding nine figures of capital.

Fathom

United States (Fathom Video Inc.) | founding date not established | about 64 million US dollars total, last a 43 million Series B in May 2025 | ADJACENT

- **Website:** https://www.fathom.ai/
- **Form:** app, on iOS, macOS and Windows, plus a bot for video calls.
- **Capture:** both. A bot-free desktop application and a dedicated iOS application for in-person meetings (https://www.fathom.ai/ios-app).
- **Time to output:** under 1 to 2 minutes from the end of the call to summary and action items in the inbox, timed independently by The RevTech Review in early 2026 on a sales call of unstated length (https://therevtechreview.com/fathom). The ingest path in that test was a video call, so the figure does not yet transfer to a room recorded through a phone microphone. It is still the fastest independently measured turnaround in this document.
- **What they do:** records the call or the room, transcribes continuously, and delivers a structured summary with action items immediately after the end.
- **Product:** a meeting bot, desktop and phone applications, and a workspace, sold per seat.
- **Tech:** software only, cloud processing, with continuous ingestion during the meeting rather than a batch pass afterwards, which is the architectural reason for the speed.
- **Who pays:** the end user or the team. The free tier gives unlimited recording and transcription and caps AI summaries at the first five calls a month; Premium and Team run 15 to 19 US dollars per user per month (https://www.fathom.ai/pricing).
- **Who subsidizes:** none. Bootstrapped to 16 million US dollars of ARR by the end of 2023 on a third-party tracker's estimate, then a 17 million Series A that reserved ten percent for its own users and took over 3.2 million from 2,148 retail investors on Wefunder at a 73 million valuation, then the Series B (https://getlatka.com/companies/fathom.video, https://growthturbine.com/blogs/25-biggest-wefunder-success-stories-updated-2025).
- **Certifications and litigation:** none published. No litigation found.
- **Traction:** over 6,500 reviews on G2 at 5.0, with claimed use by more than 300,000 companies (https://fireflies.ai/blog/fathom-ai-notetaker). That claim is carried on a competitor's blog and the company does not publish a customer count of its own.
- **Why they matter to us:** it holds the number this project's central promise has to beat. The walk-back deadline is already met by a shipping product, which makes speed table stakes rather than a differentiator, and the free tier that comes with it sets the price of the transcript at zero.
- **Their weakness:** the buyer it serves is the sales team on scheduled calls, its free tier caps exactly the two outputs that matter here, and nothing about its speed has been measured on in-person capture.

Built into the phone, or bundled into software the buyer already pays for

These are not companies a customer chooses instead of this project in the ordinary sense. They are the price floor, the distribution owner and the most likely cause of death, all at once, and one of them is also the strongest available answer to "why now". The bundled cases matter because their customer acquisition cost is zero and they never appear on a comparison page at all. The three largest Western bundlers, Microsoft Teams, Zoom and Google Meet, are in the roster at the end rather than profiled here, because each delivers its recap inside a video call rather than for a room somebody walks out of; they are the same mechanism all the same, and they are the ones the graveyard below is built from. **Apple, Google, Lark, SK Telecom.**

Apple

Cupertino, United States | founding date n/a for this purpose | n/a, public company | COMPARATOR

- **Website:** https://developer.apple.com/documentation/speech/speechanalyzer
- **Form:** platform-native, in the phone's own applications, plus frameworks any third-party application can call since iOS 26.
- **Capture:** in-person via the phone's own microphones, in Voice Memos and Notes.
- **Time to output:** no processing-time statement is published. Transcription runs while recording, so the transcript exists within seconds of stopping; the summary is a separate request.
- **What they do:** iOS 18 added on-device transcription to Voice Memos on iPhone 12 and later; Notes records, transcribes live and summarises through the on-device foundation model. For developers, iOS 26 exposes SpeechAnalyzer with a long-form SpeechTranscriber, a short-utterance DictationTranscriber and a voice-activity SpeechDetector, plus SystemLanguageModel through the Foundation Models framework for summarisation.
- **Product:** Voice Memos and Notes for the user; SpeechAnalyzer and Foundation Models for the developer.
- **Tech:** on-device speech recognition and an on-device language model, no network round trip and no per-hour bill. The published accuracy is condition-dependent and the condition is the whole story: an independent benchmark over 5,559 utterances measured SpeechAnalyzer at 2.12 percent WER on the clear read-aloud subset of 2,620 samples, against Whisper Small at 3.74 and Whisper Large V3 Turbo at 3.01 (https://gigazine.net/gsc_news/en/20260714-apple-speech-analyzer-benchmark/); on earnings22, about twelve hours of real corporate earnings-call conversation, the same framework scores 14.0 percent and WhisperKit Small 12.8. No speaker diarization is exposed to third parties, and the on-device model enforces a hard 4,096-token limit per session covering prompt, transcript and output together (https://developer.apple.com/documentation/foundationmodels/managing-the-context-window).
- **Who pays:** nobody. It is free with the phone.
- **Who subsidizes:** n/a.
- **Certifications and litigation:** n/a. Its own documentation is also the constraint every application in this document lives under: an app may be suspended following an interruption, accepting a phone call suspends it, and a backgrounded application that tries to restart capture receives a specific refusal from the privacy subsystem (https://developer.apple.com/library/archive/documentation/Audio/Conceptual/AudioSessionProgrammingGuide/HandlingAudioInterruptions/HandlingAudioInterruptions.html).
- **Traction:** every iPhone 12 and later for transcription; iOS 26 and later for the developer frameworks.
- **Why they matter to us:** two edges and both are sharp. The frameworks are the only dated reason this is buildable now that survives scrutiny, because they remove the per-recorded-hour cost and shorten the walk-back clock in the same release. They are also the argument that the phone can do this without hardware, made free, by the company that makes the phone. The why-now and the absorption risk are one fact with two signs.
- **Their weakness:** the built-in applications are a recorder with a summary button rather than a meeting product: no action items as a first-class output, no speaker labels in Voice Memos, and no interest in what happens after the meeting. Apple has historically shipped the capability and left the product to others.

Google

Mountain View, United States | founding date n/a for this purpose | n/a, public company | COMPARATOR

- **Website:** https://developers.google.com/ml-kit/genai/speech-recognition/android
- **Form:** platform-native on Pixel devices, plus frameworks any third-party application can call.
- **Capture:** in-person via the phone's own microphones, in the built-in Recorder application.
- **Time to output:** no processing-time statement is published. Transcription is real time, and an on-device summary is reported at 5 to 15 seconds once requested, at an unstated recording length.
- **What they do:** Pixel Recorder records, transcribes in real time entirely offline, labels speakers, and generates a structural summary with on-device Gemini Nano. It is free, carries no artificial length limit on raw audio beyond device storage, and has no subscription. Speaker labels have shipped

since application version 4.2 (https://9to5google.com/2022/12/07/pixel-recorder-speaker-labels/), which is the capability neither platform exposes to a third-party developer.
- **Product:** the Recorder application for the user; ML Kit GenAI Speech Recognition and Summarization for the developer.
- **Tech:** Gemini Nano through the AICore system service, 4-bit quantized at roughly a 1 GB footprint, with a 4,096-token context window and a stricter 1,024-token per-prompt limit. The developer-facing speech API is in alpha: its basic mode runs on Android 12 and higher across 15 locales, and its generative advanced mode is hardware-gated to Pixel 10 and Pixel 11 (https://developers.google.com/ml-kit/genai).
- **Who pays:** nobody. It is free with the phone.
- **Who subsidizes:** n/a. The cost is carried as a reason to buy the hardware.
- **Certifications and litigation:** n/a for the application. On-device processing is the privacy posture.
- **Traction:** every supported Pixel device; Google reported a 24 percent engagement lift on the Recorder application after the Gemini Nano summarisation feature shipped (https://android-developers.googleblog.com/2024/08/recorder-app-on-pixel-sees-boost-in-engagement-with-gemini-nano.html).
- **Why they matter to us:** it does the whole job free, offline and with speaker labels, and it withholds the speaker labels from everyone else. The one output it is not documented as producing is action items as a distinct checklist rather than lines inside a summary, and that gap is undocumented rather than demonstrated.
- **Their weakness:** it is Pixel-only, so it does not reach most Android users, and its summarisation is bound against the same 4,096-token window that limits every on-device pipeline, roughly a third of a one-hour meeting transcript. A context window bounds one generation pass and not a product, so the recorder has to be chunking an hour-long transcript and rolling the pieces up, and how that survives six overlapping speakers is documented nowhere.

Lark

Singapore and China (ByteDance) | founding date n/a for this purpose | n/a, division of a private company | ADJACENT

- **Website:** https://www.larksuite.com/
- **Form:** app, as a feature of a bundled suite the buyer already pays for.
- **Capture:** in-person via the phone’s own microphones, recorded from the Lark mobile application.
- **Time to output:** not published.
- **What they do:** Lark Minutes records a live meeting from the mobile application and routes the transcript, the summary and the action items automatically into the chat channels and CRM bases the company already works in, with no third-party subscription and no external integration to configure (https://www.larksuite.com/hc/en-US/articles/360048487750-record-a-one-on-one-call-or-meeting).
- **Product:** a chat, document, calendar and meeting suite with recording built in.
- **Tech:** cloud, inside the suite. For physical capture the company partnered with Anker on a ten-gram wearable button that syncs offline audio into the Lark cloud with no user intervention, which is hardware treated as an accessory to software rather than as the product.
- **Who pays:** the company that already bought the suite. Individuals get 300 free transcription minutes a month; paid enterprise tiers from 6 to 39 US dollars per user per month include unlimited processing (https://www.larksuite.com/en_us/plans).
- **Who subsidizes:** the suite. Meeting capture is a retention feature, not a revenue line.
- **Certifications and litigation:** none published.
- **Traction:** not disclosed for the Minutes feature.
- **Why they matter to us:** it is the clean case of the competitor a comparison page never lists. A bundled product acquires the market at a customer acquisition cost of zero, and it wins on the one axis this project has no answer to, which is that the transcript is already where the work happens.
- **Their weakness:** it reaches only companies that have already standardised on the suite, and it is effectively absent from Western markets, which is where this project’s demand evidence comes

from.

SK Telecom

Seoul, South Korea | founding date n/a for this purpose | n/a, public company | ADJACENT

- **Website:** https://www.sktelecom.com/
- **Form:** platform-native, inside the mobile carrier's own dialer.
- **Capture:** cellular call capture from inside the native dialer, with transcription, translation and summarisation, extended into a multimodal log that links with smart glasses.
- **Time to output:** not published.
- **What they do:** the Adot assistant transcribes and summarises calls for the carrier's own subscribers with no download, no app store and no extra charge.
- **Product:** an assistant built into the phone's calling experience for subscribers of one network.
- **Tech:** carrier-side and on-device services below the app store layer.
- **Who pays:** nobody extra. It is bundled with the subscription.
- **Who subsidizes:** the carrier, as a churn-reduction feature.
- **Certifications and litigation:** none published.
- **Traction:** passed ten million subscribers by mid-2025 (https://en.sedaily.com/technology/2026/08/31/sk-telecom-adds-ai-photo-logging-to-a-app-for-smart-glasses).
- **Why they matter to us:** it is the only entity in this document that does the job the dedicated recorder's biggest buyer segment is paying for, capturing the call, and it does it by owning the layer beneath the app store. It is the shape of the threat no application can answer, because it does not compete for the download.
- **Their weakness:** one carrier in one country, and the capability travels no further than the network's own subscriber list.

The dedicated recorder the buyer is choosing today

One company defines this segment in Western markets and one company shows what the segment looks like when the subscription is removed entirely. Both sell the same three outputs; the difference is where the compute happens and therefore what the recurring bill can be. **Plaud, iFlytek.**

Plaud

Delaware and San Francisco, United States | founding date not established | about 5 to 6 million US dollars disclosed, last a 4.75 million convertible note led by Carbide Ventures in April 2025 | DIRECT COMPETITOR

- **Website:** https://www.plaud.ai/
- **Form:** hardware plus app. Since January 2026 also a software-only Mac application (https://techcrunch.com/2026/01/04/plaud-launches-a-new-ai-pin-and-a-desktop-meeting-notetaker).
- **Capture:** in-person via the device's own microphones, plus cellular call capture through a vibration conduction sensor that reads chassis vibration through the phone's back glass, which is the one job no application on either platform is permitted to do. The Mac application instead "takes notes using system audio to capture the meeting", which is call capture and not room capture.
- **Time to output:** not published. The company's own pricing page states no processing or turnaround time (https://www.plaud.ai/pages/plaud-ai-plan-pricing). A 2026 hands-on review describes notes and action items organised in "seconds" once the recording syncs, at an unstated recording length (https://www.bgr.com/2225669/plaud-note-pro-review/), and the vendor's own blog says "It only used to take 30 minutes, but now it takes seconds".
- **What they do:** a credit-card-shaped recorder that attaches magnetically to the back of a phone, records locally, then uploads to a cloud pipeline that transcribes, labels speakers and writes a structured summary.

- **Product:** the Plaud Note at 159 US dollars (1,149 RMB, about 160.90 US dollars, in China), the Note Pro at 179 to 181.90, and the wearable NotePin and NotePin S at about 174.90, each paired with a phone application and now a desktop application (https://www.plaud.ai/products/plaud-note-ai-voice-recorder).
- **Tech:** proprietary hardware, 0.117 inches thick, 64 GB of local storage, up to 30 hours of continuous recording and 60 days of standby, transcription in 112 languages with speaker labels. Reported microphone configurations are two Knowles Sisonic MEMS microphones plus the vibration sensor on the Note, four microphones with a stated 5 metre pickup range and beamforming on the Note Pro, and a single mono microphone on the NotePin. Those counts come from retailer listings, review sites and comparison blogs; no teardown of any Plaud device is published. The speech and language layer is bought in, not built (https://www.tomsguide.com/ai/plaud-note-review).
- **Who pays:** the end user, twice. Device sale up front, then a subscription: free at 300 transcription minutes a month, Pro at 17.99 US dollars a month or 8.33 annualised (99.99 a year) for 1,200 minutes, Unlimited at 29.99 or 19.99 annualised, Team at 35 per user per month (https://www.plaud.ai/pages/plaud-ai-plan-pricing). Buying the device buys 300 monthly minutes and no more.
- **Who subsidizes:** none. The company is substantially self-funded, and its start was crowdfunded: the Plaud Note campaign under the iZYREC creator account raised 1,108,181 US dollars from 7,564 backers on Kickstarter between 27 June and 16 August 2023 (https://www.kickstarter.com/projects/izyrec/plaud-note-chatgpt-empowered-ai-voice-recorder). Reports of a Tencent investment describe discussions at a valuation between 1 and 2 billion US dollars, not a completed round (https://kr-asia.com/tencents-rumored-plaud-deal-points-to-looming-ai-hardware-contest).
- **Certifications and litigation:** n/a, not a regulated device, and none named. Its terms require the user to warrant a lawful basis “under the law of every jurisdiction in which the participants in the conversation are located” (https://eu.plaud.ai/policies/terms-of-service). The cloud does not retain audio once the text is generated; text and metadata are kept until one month after account cancellation, on servers in the United States. The device shows a visual recording indicator and deliberately emits no audible warning on calls. No litigation against it was found.
- **Traction:** the figures are large and they do not reconcile. Over two million devices shipped globally by mid-2026, against “more than 1.5 million devices until now” reported in January 2026; domestic Chinese sales estimated at under 100,000 units in 2024; revenue reported two ways by the same tracker, 100 million US dollars of software-only ARR as of June 2026 and a 250 million total annualised run-rate calculated in September 2025, with a stated 500 million target for 2026 (https://sacra.com/c/plaud/, https://tracxn.com/d/companies/plaudai). Read as a magnitude, not as a series: the demand is real and it is in the hundreds of millions of dollars.
- **Why they matter to us:** they are the project’s premise, and they are also the correction to it. Their sales are the demand evidence the thesis rests on, their subscription is the price this project would have to charge anyway, and the largest single block of their buyers’ stated reasons is a job an application is structurally forbidden to do.
- **Their weakness:** they own neither the speech models nor the phone, so the moat is retail shelf presence and brand; every other hardware entrant in this document was bought by a platform owner or switched off; and the discretion argument cuts both ways, since one reviewer notes that “simply recording the conversation on your phone may actually be more discreet” (https://www.bluedothq.com/blog/plaud-alternatives).

iFlytek

Hefei, China | public company, Shenzhen exchange 002230 | n/a, publicly traded | ADJACENT

- **Website:** https://store.iflytek.com/
- **Form:** hardware plus on-device software.
- **Capture:** in-person via directional and omnidirectional microphone arrays with far-field pickup, processed offline on an internal neural processing unit, with no network round trip.
- **Time to output:** not published.
- **What they do:** dictaphone-descended recorders with a screen and an onboard speech engine,

plus electronic ink tablets that convert handwriting and speech to text.
- **Product:** the Smart Recorder line (SR302, SR502) at 140 to 299 US dollars, the AINOTE 2 at 649 to 727 and the AINOTE Air 2 at 499 to 520.
- **Tech:** proprietary hardware with the company's own speech stack running on device; transcription happens offline, which is why it appeals to enterprise and government buyers who require air-gapped data.
- **Who pays:** the end user, once. The hardware premium includes lifetime free on-device transcription and there is no subscription at all.
- **Who subsidizes:** n/a. The company reports 27.82 billion CNY (about 3.89 billion US dollars) of annual revenue across all its business lines, which funds the recorder line's economics (https://www.investing.com/equities/iflytek-a).
- **Certifications and litigation:** none published.
- **Traction:** not disclosed for the recorder line specifically.
- **Why they matter to us:** it is the only vendor in this document that has removed the recurring bill, and it did so by moving the compute onto the thing the customer bought. Any product charging monthly competes against a one-time payment for unlimited use, and that is the same argument this project would have to make about the phone, made first and made by a company with a chip budget.
- **Their weakness:** most of the comparison material behind these figures sits on a competing recorder vendor's own blog; the products have little Western retail presence; and the lifetime transcription is bound to the device it shipped on, so the buyer's upgrade path is another purchase.

The neighbour: the bot that joins the video call

Included as a neighbour and never as the same product. This company records a conferencing platform's own audio stream, which is a clean near-field signal and a meeting nobody walks out of, so neither its accuracy nor its economics transfer. It is here because it sets the category's price and turnaround expectations, and because its consent posture is the one the biometric statutes reward. **Fireflies.**

Fireflies

United States | founding date not established | about 19 million US dollars raised, from Canaan and Khosla Ventures | ADJACENT

- **Website:** https://fireflies.ai/
- **Form:** app, primarily a bot that joins meetings from a web workspace, with phone applications alongside.
- **Capture:** a bot joining a video call. A phone application exists, but the product's centre is the call.
- **Time to output:** published. "Fireflies usually generates searchable transcripts with over 90% accuracy in 15 to 20 minutes", for calls under 60 minutes, from the company's own guides (https://fireflies.ai/blog/how-to-use-fireflies-ai).
- **What they do:** joins a scheduled video call as a participant, records the platform's audio, and returns a transcript, a summary and action items into a searchable workspace wired into a CRM.
- **Product:** a meeting bot plus a web workspace, sold per seat.
- **Tech:** software only, rented speech and language models, near-field platform audio rather than room audio.
- **Who pays:** the end user or the team. A free tier limited by total stored minutes rather than a monthly allowance, and Pro at 10 to 18 US dollars a month. The two figures in circulation for that free cap, 800 minutes of total storage and 400 total storage minutes per team, do not agree, and the vendor's own pricing page states neither of them here.
- **Who subsidizes:** none found.
- **Certifications and litigation:** SOC 2 Type 2 and HIPAA with a Business Associate Agreement on enterprise tiers. The company states publicly that it has "a zero data retention policy and does not train on voice data" and that "No voice-based biometric profiles are created, stored, or maintained", though those defaults apply mainly to enterprise and healthcare tiers. Defendant in *Cruz v. Fire-*

flies.AI Corp. (December 2025), which adds Illinois Biometric Information Privacy Act voiceprint claims for non-users to the Otter consent theory and is consolidated with *Brewer*; unresolved.

- **Traction:** over 16 million users claimed and about 11 million US dollars of ARR on a third-party tracker's 2024 estimate, alongside a 1 billion US dollar valuation announced on the company's own blog in June 2025 (https://getlatka.com/companies/firefliesai, https://fireflies.ai/blog/fireflies-1-billion-valuation). Those four figures, 19 million raised, 1 billion valuation, 11 million ARR and 16 million users, do not reconcile with each other and no audited number exists.
- **Why they matter to us:** it is the counter-example that makes the in-person problem visible. The bot has a clean audio feed, an explicit consent signal in the participant list and no battery or interruption problem, and it still publishes a 15 to 20 minute wait and eight figures of revenue. It also shows what a defensible consent posture looks like: no voiceprints at all.
- **Their weakness:** it structurally cannot serve the meeting this project is about, because there is no bot to invite to a conversation in a corridor, and the visible-bot notice that is its consent mechanism is exactly the thing a growing share of buyers say they are trying to escape.

The Graveyard (Companies That Failed and Why)

The deaths in this field are mostly not commercial in the usual sense. Two of the three below kept their customers and lost the company anyway; only Humane ran out of buyers, with returns running ahead of new sales inside a year of launch. Two mechanisms did the work, and they are worth separating because only one of them is available to a company that ships no hardware.

Absorbed by the platform underneath

The wrapper dies when the platform it sits on ships the feature. The named case is the clearest one this field has, and the pattern repeats at least four more times in adjacent productivity categories. **Vowel AI.**

Vowel AI

HQ not established | founded 2018 | 17.8 million US dollars raised, including a 13.5 million Series A in September 2021 | GRAVEYARD: shut-down

- **Website:** n/a (defunct); the team went to Zapier.
- **Form:** app, on the desktop and the web, as a conferencing platform in its own right.
- **Capture:** a bot in its own video call, never a room recorded through a phone microphone.
- **Time to output:** not published.
- **What they built:** an AI video conferencing platform with real-time transcription, collaborative agendas and automatic summaries, which asked the customer to move their whole conferencing stack in order to get the notes.
- **Tech:** cloud, hosted speech recognition and language models, no hardware, no proprietary corpus.
- **Final payment state:** subscription revenue with no subsidiser. A signed venture term sheet was withdrawn at the last minute against a high burn rate, which left the company immediately insolvent.
- **Why it failed:** two causes arrived together and the order matters. Zoom and Microsoft Teams shipped free native meeting summaries, which collapsed the reason to pay for a separate product; the withdrawn term sheet then removed the runway to respond. The company announced its shutdown in July 2023, and the remaining team was acqui-hired by Zapier in March 2024 to build Zapier Central (https://otter.ai/blog/smooth-transition-moving-from-vowel-to-otter-ai-as-vowel-shuts-down).
- **The lesson for us:** the asset a wrapper holds is the layer, and when the host ships the layer there is nothing left to sell but the team. The same mechanism killed Regie.ai when CRM platforms shipped native AI writing, ended Viable when its own model supplier gave enterprises the tools directly, and retired Dragon NaturallySpeaking for Mac after Microsoft acquired Nuance.
- **Pattern match risk:** 4 out of 5. The difference is only which platform: for Vowel the host was the conferencing platform, for this project it is the operating system, and both operating systems have

already shipped recording, transcription and summarisation free.

The device outlived its servers

Both entries sold the same output this project sells, both attached the output to an object, and both left their customers holding a dead object inside twelve months. **Humane, Limitless.**

Humane

San Francisco, United States | founding date not established | about 230 million US dollars raised, peak valuation 850 million, sold for a reported 116 million | GRAVEYARD: fire-sale

- **Website:** n/a (defunct); assets held by HP.
- **Form:** hardware plus app, with its own cellular connection.
- **Capture:** in-person via the device's own microphone, worn on the chest.
- **Time to output:** not published; response latency in use was reported at two to five seconds per query.
- **What they built:** a wearable AI Pin with its own camera, microphone, laser projector and cellular connection, intended to replace the phone rather than record beside it. Launched April 2024 at 699 US dollars, later 499, with a mandatory 24 US dollar monthly cellular plan.
- **Tech:** proprietary hardware with a cloud platform (CosmOS) behind it; over 300 patents and patent applications went to the acquirer (https://m.gsmarena.com/humane_ai_pin_discontinued_hp_to_take_over_ip_and_patents_-amp-66608.php).
- **Final payment state:** device sale plus a monthly subscription including cellular service, with no subsidiser. Returns exceeded new sales by the summer of 2024 and fewer than 10,000 units shipped. HP bought the assets for a reported 116 million US dollars in February 2025 (https://gizmodo.com/humane-bricks-its-ai-pin-as-it-gets-acquired-by-hp-for-116-million-2000565528).
- **Why it failed:** it asked the buyer to carry a second connected device that did less than the phone in their pocket, and the execution made the ask worse: thermal throttling on routine queries, two to five seconds of latency, and a laser projection illegible in daylight. The ending is the instructive part. Every device stopped working at noon Pacific time on 28 February 2025, and owners whose device shipped before 15 November 2024 received no compensation at all (https://www.biometricupdate.com/202502/humane-is-bricking-ai-pins-following-116m-hp-buyout).
- **The lesson for us:** the failure was not the technology, it was asking someone to acquire, charge and carry a second object for a job the phone was already holding. That is this project's thesis stated from the other side, and it is the strongest external support the thesis has.
- **Pattern match risk:** 1 out of 5. This project proposes no hardware, so this specific death is unavailable to it. What transfers is the second-order lesson: a product whose value lives on someone else's servers can be switched off, and an application is not exempt from that.

Limitless

San Francisco, United States | founding date not established | 34.3 million US dollars raised per one company database, 15 million per another | GRAVEYARD: acquihire

- **Website:** n/a (defunct); team absorbed into Meta.
- **Form:** hardware plus app, after an earlier desktop-only application.
- **Capture:** in-person via the pendant's own microphone, worn continuously rather than started per meeting.
- **Time to output:** not published.
- **What they built:** a 99 US dollar wearable Pendant that recorded the wearer's day continuously, with a physical consent chime to alert bystanders, after the earlier Rewind desktop application that recorded and indexed everything on a computer screen.
- **Tech:** proprietary wearable hardware with a cloud transcription and retrieval pipeline; the desktop product before it was software only.

- **Final payment state:** device sale plus subscription at shutdown, with no subsidiser. Acquired by Meta in December 2025; the device was pulled from sale, the Rewind application was shut down, and service ended entirely in several countries including the European Union, Brazil, South Korea, Israel, Turkey and the United Kingdom (https://mlq.ai/news/meta-acquires-ai-wearables-startup-limitless-ending-sales-of-pendant-device/).
- **Why it failed:** it did not fail commercially in the ordinary sense; it was bought and switched off, which for a customer is indistinguishable from failure. The unit economics of manufacturing a standalone wearable for a narrow professional audience never scaled, and buyers report the pendants going inert after the acquisition. Note the direction of travel: this company started as software on a device the user already owned and added hardware, and the hardware is what the acquirer wanted.
- **The lesson for us:** in this category a good team plus a capture position is worth more to a platform owner than the business is to its own customers. An application with no hardware and no proprietary corpus brings less to that table than a hardware team did.
- **Pattern match risk:** 2 out of 5. An application cannot be bricked the way a pendant can, but the acquisition logic applies: the demonstrated good outcome in this category is being bought by whoever owns the platform, and this project would arrive at that table with the least to sell.

Successful Exits (What Worked)

Every hardware company in this document that exited, exited into a platform owner, and all three deals closed between February and December 2025; only one price was ever disclosed. The category's demonstrated outcome is not a public offering or a revenue multiple; it is acquisition by a company that already ships an assistant, and in two of three cases the product the customers had bought was ended or absorbed.

Humane | Exit: assets to HP for a reported 116 million US dollars, February 2025 | Multiple: 0.14 times the last private valuation of 850 million - **What they built:** the AI Pin, a wearable meant to replace the phone. - **Why it worked, in the narrow sense that it closed:** the patents. HP took over 300 patents and patent applications, the CosmOS platform and some staff. - **Relevance to us:** it is the floor case and it prices the thing this project has none of. When a consumer AI capture company fails, what is left to sell is intellectual property.

Limitless | Exit: acquired by Meta, December 2025, terms not disclosed | Multiple: not disclosed - **What they built:** the Pendant wearable recorder at 99 US dollars, after the Rewind desktop recorder. - **Why it worked:** for the acquirer it bought a team and a wearable-capture position; the product line was ended. - **Relevance to us:** it is the exit comp this project would be measured against, and the warning inside it is that the assets bought were the team and the hardware position, neither of which a pure application has.

Bee | Exit: acquired by Amazon, July 2025, terms not disclosed by the acquirer | Multiple: not disclosed - **What they built:** a 49.99 US dollar wearable pendant sold as a loss leader in front of a 24 US dollar monthly subscription. - **Why it worked:** the product survived the acquisition rather than being switched off, and the reading offered is that the pendants are worth more to the buyer as real-world conversational data ingest next to an assistant it already ships than as a business. One company database lists the price at 7 million US dollars; the acquirer disclosed nothing, so that figure stands alone. - **Relevance to us:** it shows what the platform owners are buying, which is a capture position adjacent to a voice assistant. A phone application is not that.

Vowel AI (team only) | Exit: acqui-hired by Zapier, March 2024, after the product shut down in July 2023 | Multiple: n/a, the company was insolvent - **What they built:** an AI conferencing platform with transcription and summaries, on 17.8 million US dollars of venture funding raised from 2018. - **Why it worked, for the acquirer:** it bought an intact team that had shipped real-time meeting intelligence, at the price of a hiring round. - **Relevance to us:** this is the only software exit in the set, and it happened after the product died. It is the realistic downside shape for a software-only entrant: the team is the asset, and it is bought at team prices.

The Failure Patterns

Pattern: platform absorption. Vowel AI died when Zoom and Microsoft Teams shipped free summaries; Regie.ai, Viable and Pencil AI died the same way when their host platforms shipped the feature; Dragon NaturallySpeaking for Mac was retired after Microsoft acquired Nuance. In this project's case the host is the operating system, and both operating systems have already shipped recording, transcription and summarisation free, with Apple also exposing the whole pipeline to third parties at no per-hour cost. *We are at risk, and the exposure is structural and unmitigated: an application whose entire value is a thin layer over models the platform gives away is the most absorbable shape in software, and "no new device" is the platform's argument too.*

Pattern: the second object. Humane and, in customer terms, Limitless. Both asked a buyer to acquire, charge and carry a device for a job the phone in the same pocket could do, and both are gone; Humane's servers went dark with no refunds for most owners. *We are not at risk from this pattern. It is the pattern the project exists to exploit, and "nothing to be bricked" is a sharper and more testable version of the project's own argument than "nothing to charge".*

Pattern: distribution owned upstream of the store. Lark routes meeting notes into the chat and CRM the company already uses; SK Telecom's Adot sits in the native dialer for over ten million subscribers and never touches an app store. A bundled competitor acquires the market at a customer acquisition cost of zero and never appears on a comparison page. *We are at risk in any market where a suite or a carrier decides to ship this, and an application has no defence except being installed first.*

Pattern: undifferentiated crowding. Voicenotes, Granola, Otter, Wave, Fathom, MeetGeek, Notta, Jamie and a dozen more offer the same three outputs from the same rented models at 9 to 19 US dollars a month, and three of them give the core interaction away. When products cannot be told apart, ranking and advertising decide, and acquisition cost climbs past what a subscriber is worth. *We are at risk, and this is the pattern the thesis has no answer to: "we removed the device" is a difference from Plaud, not from Voicenotes.*

Pattern: unit-economics inversion. Cloud transcription is metered per recorded hour at roughly 0.21 to 0.62 US dollars while subscriptions are flat, so the heaviest and most attached users are the least profitable, with breakeven near 32.5 hours a month at a 15 dollar price. The category's response is visible in the identical 300-minute free caps. *We are at partial risk, and the on-device path is the one credible escape, because it moves the marginal cost of a recorded hour to zero. That escape costs the 4,096-token context window and the speaker labels the platform reserves for itself.*

Pattern: trust death by a single lost recording. A meeting recorder is bought as insurance, and one silently lost recording ends the relationship. This is documented rather than speculative: Apple's own documentation states that an application may be suspended following an interruption and that accepting a phone call suspends it, and a backgrounded application's attempt to restart capture is refused by the privacy subsystem; on Android the handset manufacturers, not the operating system, terminate background recording after screen lock. *We are at the highest risk here of any pattern in this document, because the dedicated recorder is structurally immune to it: it has its own microphone and its own battery, and no incoming call can take either.*

Pattern: the covert-capture marketing trap. Four class actions are pending and none has resolved. The one that matters most is *Chamberlain v. Granola*, because Granola processes audio on the user's own device and was sued anyway: the complaint attacks the marketing of hidden capture, not the location of the compute. Meanwhile every vendor pushes the consent burden onto the user in its terms, which protects it in contract disputes and does nothing against non-users who signed nothing. *We are at risk in exactly the same way, and more so: a two-tap product with no bot in the room and no announcement is the fact pattern the Granola complaint describes. The compliant interaction is at minimum start, announce, stop.*

The Success Patterns

The object is the acquisition channel, not the product. Plaud's device is cash-positive at the point of sale and pays for finding the customer, which is why a company shipping over two million units has raised under 6 million US dollars. The physical purchase also anchors the subscription. Remove the 159 dollars and the mechanism that funds acquisition goes with it, and nothing in this document replaces it.

Retention and seats get funded, not accuracy. Granola went from a 250 million to a 1.5 billion dollar valuation in ten months on enterprise adoption and growth, not on a published error rate. Otter's growth engine is its Business and Enterprise tiers. Nobody in this document competes on a benchmark number, and nobody publishes one.

Zero marginal cost removes the price floor for whoever gets there first. Meetily runs entirely locally and can offer unlimited transcription indefinitely; Snaply prices team seats at 5 US dollars against a 9 to 19 dollar category; MacWhisper charges about 69 dollars once; iFlytek sells lifetime on-device transcription with the hardware. Local processing is not only a privacy posture, it is a pricing weapon, and the products holding it are mostly desktop applications rather than phone applications.

Compliance can be the distribution. Jamie sells at 24 to 47 euros a month, two to three times the category, into European buyers who cannot use a United States cloud: ISO 27001 and SOC 2 Type II, European-only servers, audio deleted the moment the transcript exists, and a contractual guarantee that customer data never trains a model. The certification is the product, and the same posture is what the biometric statutes reward.

The platform owner buys the capture position. Bee into Amazon, Limitless into Meta, Humane's patents into HP, all within twelve months. What they bought was a place next to an assistant they already ship, plus a team.

Speed is the one axis nobody has claimed and one company has already taken. Fathom delivers summary and action items in under one to two minutes on a measured video call while the category default is 10 to 20 minutes, Otter publishes 2.5 to 6 minutes to a summary for a 30-minute meeting, and not one vendor publishes a guaranteed maximum. The unclaimed position is not "fast", which is taken, but "fast with a number attached, on a room recorded through a phone".

Also Found, Not Profiled

Every company the sweep surfaced that did not earn a profile, one row each, alphabetical. This is the coverage record: it shows how wide the search was, and it lets a later pass promote a row without repeating the research. A blank or hedged cell is the correct outcome where nothing was established.

Company	What they sell	What we know
Aiko	On-device transcription for Apple Silicon	One-time purchase rather than a subscription; zero marginal cost per user
Amberscript	GDPR-compliant transcription for European academic and corporate buyers	Positioned as a replacement for human transcription services
Anker AI Recording Bean	Ten-gram wearable recording button	Built as an accessory to ByteDance's Lark; syncs offline audio into the Lark cloud with no user step
AssemblyAI	Hosted speech recognition sold to developers	List price about 0.37 US dollars per recorded hour, a cheaper tier quoted at 0.15
Audionotes	Mobile-first transcription for Indian users	Unfunded, Mumbai; built for rapid Hindi-English code-switching

Company	What they sell	What we know
Basil AI	On-device note taker on Apple's native frameworks	The named example of the posture where no vendor server ever holds the recording
Bee	49.99 dollar pendant in front of a 24 dollar monthly subscription	Acquired by Amazon July 2025 and continued; acquirer disclosed no terms, one company database lists 7 million US dollars
Bliro	European bot-free meeting assistant	Named in the European compliance cluster; no figures established
Claap	Meeting recorder writing structured notes into Salesforce and HubSpot	SOC 2 Type 2 and GDPR with data-residency options; scores deals against named sales frameworks
Coconote	Lecture and study note taker (Quizlet Inc.)	4.8 stars over 17,000 iOS reviews; no speaker labels and no action items, it produces flashcards and quizzes
CraftNote	Android note taker with background recording	Records with the screen locked; returns transcript, summary and action items as a distinct artifact
Daglo	Korean transcription focused on dialect accuracy	Competes locally with Naver and SK Telecom
Deepgram	Hosted speech recognition sold to developers	Nova-3 at 0.0043 dollars a minute batch (0.258 an hour) and 0.0077 streaming; 200 dollars of free credits on signup
DingTalk A1 (Alibaba)	Meeting capture bundled into an enterprise suite	1,000 free minutes a month; professional tier 599 RMB a year, premium 1,499
Dragon NaturallySpeaking	Desktop dictation software (Nuance)	Mac version discontinued 2018, Dragon Home 2023, after Microsoft acquired Nuance; the consumer line ended with the acquisition
ElevenLabs	Hosted speech recognition (Scribe)	Quoted at 0.22 to 0.48 US dollars per recorded hour
Even Realities	Display smart glasses	150 million US dollar pre-Series B led by Meituan and Tencent at a 1 billion valuation, July 2026; a completed round, unlike the Plaud reports it is easily confused with
Fellow	AI meeting notepad for venture and management teams	About 41.2 million US dollars raised; revenue not disclosed
FinMate AI	Meeting capture for financial advisers	Syncs redacted notes into wealth-management systems such as Redtail and Wealthbox
Friend	AI pendant companion	2.5 to 3 million dollar seed (Pace Capital, a16z, Caffeinated Capital); 1.8 million spent on the domain name; about 348,000 dollars of revenue from 3,000 units
Gong	Revenue intelligence over sales calls	Holds a proprietary graph over business-to-business calls; no outside party has obtained access to it
Google Meet	Meeting notes generated by the platform that hosts the call	"Take notes for me" writes a Google Doc carrying a summary, a decisions list and "Next steps", described as "actionable items and assigned tasks"; gated on "an eligible Google Workspace edition or Google AI plan" and the help page names no edition
Groq	Hosted speech recognition on its own silicon	Whisper-v3 quoted at about 0.04 US dollars per recorded hour, the cheapest found
Hedy	Real-time conversation coach, on device	A wholly free mode on local models with no minute cap; Pro at 12.99 to 14.99 a month

Company	What they sell	What we know
HiDock H1	Desk speakerphone dock with transcription	554,444 US dollars from 2,431 backers on Kickstarter and Indiegogo, October to December 2023; about 250 dollars with free transcription in its own software
Jamie	Bot-free meeting assistant for European buyers (Germany)	ISO 27001 and SOC 2 Type II, European-only servers, audio deleted once the transcript exists, contractual no-training guarantee; 24 to 47 euros a month
Leexi	European meeting assistant	Named in the European compliance cluster; no figures established
MacWhisper	On-device transcription for macOS	Pro at a one-time fee of about 69 US dollars, no subscription
Meetily	Fully local meeting assistant, open source	Zero cloud cost per user and unlimited transcription; free for individuals, 120 to 360 dollars a year for managed team deployments; desktop, not phone
MeetGeek	In-person and virtual note taker on iOS and Android	About 2.56 million US dollars raised; 4.9 stars from 157 iOS reviews; measured at about 17 minutes to output on a 33-minute recording; 9.99 to 19 a month
Microsoft Teams	Meeting recap bundled into the suite the buyer already licenses	Copilot “summarizes key discussion points, including who spoke and what they said. It also suggests action items”; intelligent recap ships in the Teams Premium add-on licence or the Microsoft Copilot licence, and live transcription must be on for either
Mobvoi (TicNote)	Wearable AI recorder, listed in Hong Kong	578 RMB a year for 1,500 minutes, 1,188 for 6,000
Naver Clova Note	Korean transcription on the HyperCLOVA X model	Tuned for Korean honorifics and dialects; SK Telecom’s main domestic rival
Noota	European meeting assistant	Named in the European compliance cluster; no figures established
NoteGen	AI transcription application (BUD Technologies)	Independently timed at 2 to 3 minutes for transcription plus action-item extraction on a 30-minute file
Notta	Transcription for international business teams (Japan)	31.8 million US dollars raised (Granite Asia, GL Ventures, Mizuho); claims over 8 million users; free tier 120 minutes a month with a 3-minute per-file cap; claims 98.86 percent accuracy naming no test set; also sells a 149 dollar Notta Memo recorder
Omi	Open-source wearable recorder	The destination users named when leaving Limitless after the Meta acquisition
OpenAI	Hosted speech recognition and language models	File transcription at about 0.006 US dollars a minute; also the model layer reviewers say the recorders run on
Pencil AI	Generative creative tool	Acqui-hired in 2024 after Meta and Adobe shipped the same features free
Read AI	Meeting, email and message intelligence	81 million raised; 50 million Series B led by Smash Capital at a 450 million valuation, October 2024; adding 100,000 accounts a week; valued on Fortune 500 penetration, not consumers
Regie.ai	AI sales copy tool	Shut down in 2024 after CRM platforms shipped native AI writing free
Rev	Human and AI transcription service	AI transcription at 0.25 US dollars a minute; Rev Max at 29.99 a month; no perpetual free tier

Company	What they sell	What we know
Rewind	Desktop recorder that indexed everything on screen	The same company as Limitless; shut down in December 2025 after the Meta acquisition
Samsung Voice Recorder	Built-in recorder on Galaxy devices	Reported to offer transcription and general summarisation without distinct action-item extraction; the report carries no citation and is unverified
Senstone	Wearable dictation pendant	About 100,000 US dollars from roughly 1,000 Kickstarter backers, February 2017; shipped, then faded
Snaply	Local-first meeting notes	No default cloud note store; free individual tier and team seats at 5 US dollars against a 9 to 19 dollar category
Sourcenext AutoMemo	Japanese dictaphone with transcription	Priced under the 100,000 yen threshold that lets a Japanese business expense it inside the same fiscal year
Speakwise	iOS note taker for named professional roles	59.99 US dollars a year with no free tier; claims 95 percent or better accuracy naming no test set; surfaced mostly through its own content marketing
Superwhisper	On-device transcription (macOS, Windows, iOS)	Local models; free tier plus about 5.99 US dollars a month; claims hundreds of thousands of daily active users
tl;dv	Meeting recorder with a mobile in-person notetaker	10,000-plus Play Store downloads for the mobile application; 29 to 65 US dollars per seat; a widely repeated Series A figure traces to an unrelated finance-industry primer and does not survive checking
Tinrec	Mobile transcription with an AI chat query, Asia-Pacific	Instructs users to put calls on speakerphone to work around the call-recording block; free or nominal local tiers; sourced to its own blog
Umevo	AI voice recorder card	149 to 169 US dollars with unlimited transcription in year one, dropping to 400 minutes a month in year two; much of the category's published hardware pricing traces to this competing vendor's own blog
Viable	Customer-feedback analysis	Silent shutdown after its own model supplier gave enterprises the same capability directly
Zocks	Meeting capture for wealth management	Writes secure redacted notes straight into industry-specific systems
Zoom	Meeting summaries inside the conferencing platform the buyer already pays for	The vendor's own assistant page describes a basic tier included with the paid plan at no extra cost, capped at meeting summaries for three hosted meetings a month; what unlimited summarisation costs is not stated there

Watchlist

Dated 2026-09-01. A fired trigger means this document is owed a refresh.

What to watch	Who	What it would change	Where it shows up
Any independent user, download or revenue figure appearing for the free product that does this exact job	Voicenotes	Decides whether the free application is losing on distribution or on capability, which decides whether this category is served or merely undiscovered	Company database profiles, app-store rank trackers, the company's own posts
A ruling on the motions to dismiss in the pending note-taker class actions	Otter, Fireflies, Granola, Microsoft	Sets whether silent device-level capture without all-party consent is defensible as marketed, which decides the shape of the interaction	Northern District of California and Western District of Washington dockets
Either platform exposing speaker diarization to third-party developers	Apple, Google	Removes the one capability the platform reserves for its own recorder and changes what any application can claim	WWDC session list, iOS release notes, ML Kit release notes
Phone-only capture shipping from the hardware incumbent on iOS or Android	Plaud	Removes this project's differentiator from inside the incumbent, which already ships a software-only desktop application	plaud.ai product pages, app store listings
In-person mobile capture bundled into a Western suite or carrier	Microsoft, Google, Zoom	Puts a zero-acquisition-cost competitor in the buyer's existing subscription, the pattern that already closed the Asian market	Product release notes, enterprise product roadmaps
Any vendor publishing a guaranteed maximum time from the end of a recording to a finished summary	any competitor	Turns the walk-back promise from an open claim into a number that has to be beaten	Vendor pricing, help-centre and status pages
Any vendor publishing a summary or action-item quality figure with a named test set	any competitor	Makes output quality a decidable purchase axis for the first time in this category	Vendor engineering blogs, conference talks
A change to background audio capture or interruption behaviour on either platform	Apple, Google	Moves the highest technical risk in this document in either direction	Platform release notes, developer forums
Another wearable recorder acquired by a platform owner	Meta, Amazon, Google	Confirms the exit shape and reprices what a software-only company is worth in it	Acquisition announcements