

Recapp

Brainstorm, the adjacent ideas

The ideas a panel of experts generated for this project, each one scored and each one audited against the evidence this project has actually gathered.

2026-09-02



Prepared by Copundit

AI for Advanced Ideation™

The ideas a panel of experts generated for this project, each one scored and each one audited against the evidence this project has actually gathered. Seven specialists were given the same question and the same approved evidence base, proposed independently without seeing each other, and then scored every surviving idea on four measures. A separate reader, who proposed nothing of their own, checked each idea against the evidence and recorded exactly where it rests on something nobody has measured. Ideas that contradict this project's own rules were kept and argued rather than smoothed away, and every idea that was folded into another one is listed at the end, so nothing generated here is hidden.

Contents

The question put to the panel	4
How to read the scores	4
What the audit found before anything was scored	4
The same product, better	5
S2 (folds M1) , The recording-is-alive watchdog	5
P5 , The recording that tells you it is still recording	5
S3 , Publish the entity error rate, then move it	6
M3 , Seat-position speaker labels from the ambisonic channel	7
M6 (folds T3) , A deadline that can be published	7
S1 , Ambisonic front end on the iOS 26 capture path	8
V1 (folds P3) , The announcement inside the recording	9
V2 , Nothing outlives the summary	9
P4 , Record before install, pay off-store	10
G2 , Unmetered capture, priced per finished meeting	11
T5 , The remorse window	11
T1 , Sell readiness, meter the hours	12
T2 , One receipt a year, issued by us	12
An adjacent product for the same buyer	13
S5 (folds V3) , Action items that abstain, with the audio behind them	13
G4 (folds T4) , The recap the room receives	14
A different buyer for the same mechanism	14
M2 , Crash reporting for the microphone	15
V6 , The room where the announcement is already free	15
S4 + M5 , The phone-condition recording programme and the rig that makes the corpus possible	16
P1 , Pre-installed on the fleet Google does not reach	17
V4 , The works council package	17
G5 , The visit note that lands in the record	18
The pivot tier	19
M4 , The meeting is the array	19
T6 , Free at the person, paid at the company	20
V5 , The record you speak yourself	20
P2 , Call recap licensed into the carrier's dialer	21
S6 , Phone-condition recognition sold by the hour	21
G6 (folds F6) , The accessory we do not build	22
What was dropped, and why	23

What this slate does not answer

23

The question put to the panel

The evidence floor is approved and the pitch has been re-scored against it. Two of its five claims are dead: the phone's microphones are not reachable by an application as an array, and the person buying the 159 dollar recorder gives reasons an application cannot supply. What adjacent ideas does this evidence open? Propose ideas that stay inside this project's scope constraints where you can, and argue a pivot openly where the evidence points outside them. For each idea: what it is in one paragraph, which use-case bucket it sits in, who pays for it, and what would kill it.

How to read the scores

Novelty is how far the idea sits from what anyone in this market is already selling. Defensibility is how hard it would be for a competitor, or for the operating system underneath the product, to copy the idea or simply absorb it. Time-to-market is how soon it could be in a user's hands given what the evidence says is already possible today. Strategic fit is how closely it serves the customer and the promise this project was started on. Every total is the sum of four averages taken over seven independent specialists who scored the same slate without seeing each other's marks; the averages were counted mechanically, and the order inside each table is by total and has not been adjusted.

What the audit found before anything was scored

Nothing on this slate came back green. Green would mean an idea whose mechanism, whose payer and whose kill test all rest on evidence that has already been established, and the auditor's finding is that such an idea is not currently constructible in this project: the data inventory returns nothing on all five of the claims the pitch needs data for, and every company figure in this field carries a footnote while every customer figure is borrowed from a neighbouring market. Forty of the forty-one ideas came back yellow, which means the idea stands but rests on at least one thing nobody has measured, and in most cases the seat that proposed the idea named that gap first, which the auditor calls the round's strongest quality signal. Red means the idea contradicts something this project has already investigated and found false.

One idea was red. A founder-seat proposal to have the phone speak the recording notice out loud claimed that the two-tap interaction survived intact because the announcement is fired by the start tap and is therefore not a third one. The competitive map and the evidence base both state, in the same words, that the compliant interaction is at minimum start, announce, stop, and the consent-law work refuses that exact re-labelling in advance: the announcement is a spoken act in a room full of colleagues, so the cost is the utterance, not the finger on the glass. The idea relocated the cost to the machine and then declared it gone, which makes it an undeclared contradiction of this project's own two-taps rule. The mechanism itself is not dead, and three other seats propose it while pricing the cost honestly.

Underneath the slate the auditor found two shared assumptions rather than forty-one independent ideas. Six ideas from six different seats, with a seventh losing half its argument alongside them, all rest on one unverified proposition: that a recording fails silently often enough, on the handsets this customer actually holds, to be worth buying insurance against, and that the failure is detectable from inside a running application. Both circulating failure rates for that have been struck from this project's record as unsourced, no shipped recorder has published a defect rate against which any target could be judged, and the specific case where the microphone keeps delivering buffers full of digital silence may not be detectable from inside the application at all. The second shared assumption sits under a different seven ideas: each takes a percentage from a sample of sixty-four stated purchase reasons as a statement about demand, when that sample was assembled from forums of people who already own the hardware, so it measures why buyers stayed with a device rather than what a buyer still choosing would do. Every percentage on the slate inherits that selection. Read together, the slate is closer to two bets than to forty-one ideas, and both bets are one measurement away from resolution in either direction.

The same product, better

ID	Novelty	Defensibility	Time-to-market	Strategic fit	Total
S2 (folds M1)	3.9	2.3	4.4	4.9	15.5
P5	4.0	3.0	3.7	4.7	15.4
S3	4.1	2.0	4.9	3.9	14.9
M3	4.6	3.0	2.9	4.3	14.8
M6 (folds T3)	3.9	2.4	3.0	5.0	14.3
S1	5.0	2.3	2.7	4.0	14.0
V1 (folds P3)	4.0	2.7	4.1	3.1	13.9
V2	2.7	3.0	4.4	3.7	13.8
P4	4.3	1.7	3.0	4.7	13.7
G2	4.0	2.1	3.0	4.4	13.5
T5	2.6	1.1	5.0	4.7	13.4
T1	3.3	1.1	4.7	3.3	12.4
T2	2.3	1.0	5.0	3.6	11.9

S2 (folds M1) , The recording-is-alive watchdog

- **Proposed by:** the speech and voice machine learning seat.
- **Who pays:** the individual meeting-goer, on a consumer subscription bought inside the app store, in the paid consumer tier rather than behind an enterprise security review.
- **What kills it:** on one handset, force both failure states, an accepted incoming call mid-recording and the phone closed in a bag, and compare the per-frame audio statistics against a genuinely quiet room recorded on the same handset; if the failure buffers are not separable from live-but-quiet audio, the detector cannot exist.
- **Audit verdict:** Yellow, the mechanism stands but the detectability question is exactly the thing nobody has answered; marked a duplicate of the capture-survival group and kept as that group's technical specification.

The evidence base names the silently lost recording session as the only mechanism in this round that is both buildable and unclaimed, and no capture-defect rate has ever been published for any shipping recorder on either platform. The self-check runs on three signals the speech field already uses: runs of exactly-zero audio samples, which a live microphone carrying its own dither and room noise floor cannot produce; a long-window noise-floor and spectral-flatness track, because a microphone that is live but muffled in a bag or face down loses high-frequency energy and gains cloth-rustle transients; and a speech-presence rate from an always-on voice activity detector of the Silero class, about 2 MB in size. What reaches the user is a warning rather than another tap, so the two-tap interaction holds, and the detector is specified as false alarms per recorded hour at a fixed miss rate on true failures rather than as an accuracy, because the quiet room that is recording perfectly well dominates deployment time. The technology survey already fixes the study shape, at least 500 sessions from at least 50 real users across both platforms with a pass above 99.0 percent defect-free completion, and notes there is no category baseline to judge that bar against; folded into this entry is a pre-flight probe that refuses to accept a meeting until it has proved, on this handset and this operating-system build, that a session survives a screen lock, an application switch and a simulated interruption, and that reports the exact setting that fixes it rather than reprinting the community sites' multi-step instructions.

P5 , The recording that tells you it is still recording

- **Proposed by:** the consumer platform strategist.

- **Who pays:** the individual consumer, buying away from the app store, because inside it the store is the merchant of record and a seller who cannot issue its own refund cannot make its own guarantee; team buyers sit above that on the directly invoiced rail.
- **What kills it:** test whether a backgrounded recorder can reach a user through a phone that is locked, silenced and face down, the exact state this product lives in; if it cannot, the guarantee is unenforceable.
- **Audit verdict:** Yellow, resting on the same unanswered detectability question; marked a duplicate of the capture-survival group and kept as that group's commercial frame.

The competitive map rates trust death by a single lost recording as this project's highest risk of any pattern in the document, and the technology survey supplies the failure modes: on iOS the privacy subsystem refuses a backgrounded application's attempt to restart the microphone after a call, returning OSStatus 561145187; developers report an input tap that keeps firing while delivering zero-filled buffers, so the application writes a valid-looking file of digital silence; and on Android four named handset manufacturers kill the recording service regardless of correct interface use. The evidence base records that nobody has measured capture survival on any handset, that both failure rates circulating in the field are unsourced, and that no vendor publishes a guaranteed maximum time to a finished summary, so the two things a buyer of insurance wants, does it survive and when does it finish, are both unclaimed. The product continuously verifies that audio is really arriving, escalates to the user the moment it is not, writes through to storage in short increments so a hard kill costs seconds rather than the session, and publishes both of the numbers it measures. A platform owner will never occupy this position, because it would have to ship a feature whose whole content is a warning that its own operating system may take the microphone away, and it will never publish a survival rate for its own audio stack.

- **A seat objected:** the merchant-of-record argument is right about invoicing but overstated as a bar on the guarantee itself, since the store does issue refunds and a developer can honour a remedy out of band, so the real reason to sell away from the store here is owning the renewal, not the ability to promise.
- **A seat objected:** that same argument is Apple-only, because a Google Play developer can refund an order from the developer console, so the guarantee does not force the off-store route on Android and the arithmetic differs per store.
- **A seat objected:** the escalation path is the load-bearing part and the entry never names its gate, since on iOS a silenced or Focus-suppressed phone drops ordinary notifications and the only documented route through is the Critical Alerts entitlement, which Apple grants on request and a recorder cannot assume it will keep.

S3 , Publish the entity error rate, then move it

- **Proposed by:** the speech and voice machine learning seat.
- **Who pays:** the individual meeting-goer on a consumer subscription inside the app store, against a competitive map that records the unclaimed position as not speed, which Fathom has taken, but fast with a number attached on a room recorded through a phone.
- **What kills it:** mark named-entity spans on the reference transcripts of one existing far-field meeting corpus and score entity errors separately from overall word errors on the same decode; if entity error tracks overall word error within a small margin there is no claim to make, and settling it needs no new recording.
- **Audit verdict:** Yellow, and the cheapest test on the slate that needs no new audio; no duplicate and no pivot, with the gap being that this project has not confirmed the licence on one of the corpora the scoring would run on, and that nobody in this category has ever been paid for an accuracy number.

No named-entity word error rate on far-field single-channel meeting audio exists in any language on any corpus, and everything the pitch leans on downstream is conditional on that missing number. The finding that summarisation tolerates transcription errors, where systems above 50 percent time-constrained speaker-attributed word error rate produced summaries scoring roughly on par with systems near 11 percent (arXiv 2507.18161), holds only when the errors fall on fillers, false starts and conjunctions, while

far-field room acoustics smear consonants so that proper nouns, numbers and acronyms fail first, and those are precisely the words an action item is made of: who, by when, which account. Reporting the number costs no new audio, because the AMI and AliMeeting meeting corpora both ship word-level verbatim transcripts alongside a close-microphone reference, so entity spans can be marked on the reference and scored against the far-field decode. The product lever that moves the number is already shipping rather than research, contextual biasing on a user's personal entities in the PROCTER line of work at a reported 44 percent relative improvement overall and 57 percent on rare personalised entities, carried here as a direction and never as a number because it states no audio condition and no test set, and the cost is that a privacy-positioned product starts reading contacts and calendar.

M3 , Seat-position speaker labels from the ambisonic channel

- **Proposed by:** the mobile audio platform engineer.
- **Who pays:** the individual buyer inside the app store; the same posture opens the team and enterprise seats invoiced direct, because a label derived from where a sound came from never builds a voice vector, which is what the Illinois Biometric Information Privacy Act names as a biometric identifier at 740 ILCS 14/15(b).
- **What kills it:** make one four-component spatial recording of three people seated at a table and check whether the horizontal components separate their directions at all; if the platform has already collapsed the spatial field into a single steered channel before handing it over, the idea ends on that one file.
- **Audit verdict:** Yellow, with two gaps beyond the shared one, that the stable-seat premise appears in no evidence this project holds and that the path exists on one platform only; a near-duplicate of the other ambisonic entry, and both were kept because they live or die on the same single recording.

The technology survey records that iOS 26 adds capture of First Order Ambisonics (FOA), a four-component recording that encodes the direction each sound arrived from, through the standard media-writing interface, and read carefully this does not contradict the platform finding that killed the microphone-array claim, because it is not raw per-microphone audio: the phone computes the encoding itself and what it releases is a spatial description rather than the individual capsule signals. In a meeting people keep their seats, so a direction is a stable per-speaker key for the length of a session, and clustering utterances by direction is a route to the one capability the platforms withhold, since Google's own recorder has labelled speakers since its version 4.2 and neither platform exposes speaker separation to anyone else. Two engineering conditions travel with it: the audio session must not declare a voice-communication mode, because the system then inserts voice-processing units whose automatic gain control and noise suppression destroy the spatial information, and the path is iPhone-only on iOS 26, so the Android product gets speaker labels some other way or does without. The technology survey calls this measurement the highest-value trigger in the document and states that nobody has made it in either direction.

- **A seat objected:** a direction is a stable per-speaker key only if the phone does not move, and the term that actually rotates every key is the handset's own orientation rather than where people sit, while direction cannot separate two talkers on the same bearing, which on a long table is the common case, so the test should seat three people at distinct angles and then repeat it with two of them in line.

M6 (folds T3) , A deadline that can be published

- **Proposed by:** the mobile audio platform engineer.
- **Who pays:** the individual buyer inside the app store, the rail the whole category's consumer tier sits on.
- **What kills it:** time one real ninety-minute meeting end to end on the newest available phone with the recognizer running during the meeting, and see whether the remaining tail plus the summarisation pass lands inside the walk back after a session of thermal load; if it does not fit on the fastest handset it fits on nothing.

- **Audit verdict:** Yellow, the strongest engineering statement of the published-latency group and the only one that carries its own cost; marked a duplicate of that group, which was folded into it.

The technology survey records that the latency intuition is inverted: data-centre silicon is roughly forty to sixty times faster at raw inference than a recent phone, and yet cloud batch products measure slower end to end, because queueing and transport rather than computation are what the user waits for. If recognition is finished at the moment the user taps stop, the wall-clock wait stops being a function of a queue and becomes a function of transcript length, which is bounded, measurable and therefore publishable, and the evidence base fixes what is unclaimed here: several vendors publish typical processing times and none publishes a guaranteed maximum. The cost is carried with the idea rather than around it, since the phone's on-device language model holds 4,096 tokens of context against an hour-long transcript of eleven to thirteen thousand, which makes the post-tap pass a chunk-and-merge whose quality nobody has evaluated. Folded in from the go-to-market seat is the move that turns an engineering property into a reason to switch at the moment of purchase, putting the number into the terms of sale with a remedy behind it rather than into the marketing copy, sold away from the app store because inside it the store is the merchant of record; and the capture-integrity dividend only this seat noticed is that a session that looks healthy while producing no words for a sustained window is either a silent room or a revoked microphone, and the product can ask which.

- **A seat objected:** the refund argument folded in here treats the store being the merchant of record as a platform fact, but that is Apple only, since a Google Play developer can refund an order from the developer console, so the arithmetic has to be done per store.

S1, Ambisonic front end on the iOS 26 capture path

- **Proposed by:** the speech and voice machine learning seat.
- **Who pays:** the individual meeting-goer on a consumer subscription inside the app store, which the funding map records as the category default for this shape of product.
- **What kills it:** one recording of two talkers at known, well-separated angles through the iOS 26 spatial path on a single handset; if the per-frequency direction estimates collapse to a single mode, or the three gradient components turn out to be a fixed linear function of the pressure component, the phone is still a one-channel device.
- **Audit verdict:** Yellow, the cleanest mechanism argument in the round, with the gap on the payer rather than on the physics, because what it produces is accuracy and no buyer in this category has ever been recorded paying for accuracy; a near-duplicate of the seat-position entry, same measurement, different product off it.

First Order Ambisonics (FOA) is a four-component description of a sound field, one omnidirectional pressure channel plus three pressure gradients along three perpendicular axes, and iOS 26 added FOA capture through the standard media-writing interface, which is the one route to a spatial front end that does not need the raw per-microphone access both platforms withhold. It matters because guided source separation, the 2018 statistical front end that every top system in two consecutive CHiME challenge rounds still uses, fits a spatial covariance matrix from the phase difference between microphones and is mathematically undefined on a single channel, whereas from the pressure and gradient channels an application can compute a direction-of-arrival vector for every time-frequency bin, and the parametric spatial audio line of work builds beams and separation masks from that four-component format alone. The size of the prize is the one measurement the field has made on identical audio, 22.2 percent single-channel time-constrained speaker-attributed word error rate against 10.8 percent with a microphone array on the 170 blind NOTSOFAR-1 evaluation meetings (arXiv 2501.17304), which is one system's score on one evaluation set and not a bound on anything. Two conditions travel with it, never activating a voice-communication audio session because the system then inserts voice-processing units whose automatic gain control and noise suppression destroy the phase the direction estimate reads, and verifying the sample rate actually delivered, and the honest exposure is that the path is iPhone-only and gated on iOS 26.

V1 (folds P3) , The announcement inside the recording

- **Proposed by:** the recording-consent and privacy counsel.
- **Who pays:** the individual meeting-goer on a consumer subscription inside the app store, the category default and the rail every shipping application in this field already sits on.
- **What kills it:** pull the statutory text of the twelve in-person all-party-consent jurisdictions; if those statutes require consent before capture rather than evidenced within it, the artifact buys lawfulness in Washington alone and the third interaction was paid for nothing, and settling it is a reading exercise rather than a build.
- **Audit verdict:** Yellow, and a pivot the seat declared in its kill test rather than in its heading, because the spoken announcement is a third interaction charged against this project's two-taps rule; marked a duplicate of the announcement group and kept as its strongest statement.

On the start tap the application opens capture and shows a one-line script for the user to say out loud, then listens for that script inside its own captured audio and pins the detected utterance as the first anchored segment of the transcript, writing the time, the detected wording and the recording identifier into a consent record bound to the file. The consent-law research reports Washington as the one jurisdiction where a captured announcement is the compliance mechanism itself rather than proof of it; everywhere else it is the evidence no defendant in this category holds, and its absence is the whole complaint, since the competitive map states that a two-tap product with no bot in the room and no announcement is exactly the *Chamberlain v. Granola* fact pattern, brought against a product that was already processing on the user's own device. The binding in-person all-party set is twelve jurisdictions, thirteen on the cautious Missouri reading, with Nevada out and Oregon and Hawaii in, and the strictest-standard rule attributed to *Kearney v. Salomon Smith Barney* makes all-party consent the national operating default, so the move is to stop treating the announcement as a legal chore bolted onto the interaction and sell it as the deliverable, the only recorder that hands you a record showing you told the room. Folded in is the absorption argument: a platform owner structurally cannot ship this, because ambient recording that announces itself, logs who was told and keeps that record next to the audio puts the shipper inside the wiretap liability chain for its entire install base, which is what *Basich v. Microsoft* is, so Apple and Google ship a recording indicator and stop there.

- **A seat objected:** spotting a known announcement phrase inside your own capture is an easy keyword task but it is not free, and the operating point that matters is the false reject in a reverberant room, because telling a user who did announce that they did not is worse than shipping no consent record at all.

V2 , Nothing outlives the summary

- **Proposed by:** the recording-consent and privacy counsel.
- **Who pays:** the individual, on a one-time lifetime purchase, which the funding map records as live only where marginal cost is near zero and where every instance in this field runs the model on the buyer's own hardware, with a direct subscription sold away from the app store as the fallback if the on-device path does not hold.
- **What kills it:** read the verified-purchaser reviews on the retail listings coded for one question only, does the buyer go back and re-listen to the audio; if replaying the recording is a stated use, destroying the audio at output removes a function the buyer thinks they paid for.
- **Audit verdict:** Yellow, with the mechanism entirely inside the product's own control, which is rare here; the gaps are that the lifetime rail depends on the same unmeasured on-device path as several other entries and that the posture is not unclaimed, since Jamie already deletes audio at transcript time and is paid for its certifications rather than for the deletion.

When the three outputs are produced the audio is destroyed, and the speaker embeddings, the numerical voice vectors the speaker-separation step produced in order to tell one talker from another, are destroyed with it, so no vector ever crosses a recording boundary and no enrolment store exists anywhere in the system. The technology survey draws the legal line in exactly this place: anonymous ephemeral clustering inside a single recording is the described safe harbour, while tying a voice vector

to an identity by enrolment, a calendar invitation or an email address is the transition into named biometric identification, which is the specific allegation in *Cruz v. Fireflies.AI Corp.* and *Basich v. Microsoft Corp.* under the Illinois Biometric Information Privacy Act. The commercial precedent is compliance as distribution: Jamie holds ISO 27001 and SOC 2 Type II certification, keeps European-only servers, deletes audio the moment the transcript exists, contractually bars model training, and sells at two to three times the category price on that posture. Nothing is given up by deleting, since selling or licensing the recorded audio is recorded as a structurally closed way to make money here, and the General Data Protection Regulation (GDPR) Article 17 duty to isolate and permanently delete one named individual's voice and contributions on demand is trivially satisfied by a system that kept neither.

- **A seat objected:** a one-time lifetime price is where consumer software in this field goes to die, because support, model updates and operating-system churn keep costing after the revenue stops, and this project's own evidence base has already struck out the single lifetime-tier precedent the rail leans on.

P4 , Record before install, pay off-store

- **Proposed by:** the consumer platform strategist.
- **Who pays:** the individual consumer on a direct subscription sold away from the app store rather than inside it, where the funding map's worked example on a 15 US dollar subscription nets 14.26 through a United States web checkout against 12.75 through in-app purchase at the small-business rate, and the point is that the billing relationship, the renewal and the refund all stay with the seller.
- **What kills it:** read the platform policy for the zero-install surfaces, Apple's App Clips and Android Instant Apps; if microphone capture or the background audio mode is unavailable to them, the idea dies in one paragraph of published policy and no build is needed.
- **Audit verdict:** Yellow, with fully sourced commission arithmetic and a position the evidence base explicitly names as unclaimed; the gap is that this project holds no platform policy text of any kind, so the whole idea rests on one unread document, which is also why its kill test is among the cheapest here.

The survey of the existing application shelf counts a seven-step first-run gauntlet for any third-party application, launch, account creation, biometric prompt, microphone permission, notification permission, dismissing the subscription upsell, then record, against one pinch on the dedicated device, and finds that returning users already reach one to two taps on Granola through a watch complication, a Live Activity, the Dynamic Island or an Android Quick Settings tile, and on Pixel Recorder through a widget, concluding that two taps is table stakes for a returning user, unreachable for a new one, and that the first-run path is the one genuinely unclaimed position on the shelf. Every surface that compresses that path is owned by the platform and handed to developers at no charge, and this category uses those surfaces only for people who have already installed. The shape is that the first recording starts from a link, a code or a system action with no install and no account, and the account is asked for at the moment the three outputs are handed over, which is also the moment the buyer has seen the product work. That inverts the crowding pattern the competitive map describes, where store ranking and advertising decide and acquisition cost climbs past what a subscriber is worth, and it puts the purchase on the seller's own checkout so that one store review cannot close every route at once.

- **A seat objected:** the zero-install surfaces do not need a build to settle either way, because App Clips are documented as not permitted to run background audio and Google has announced the retirement of Android Instant Apps, so the surface that would carry a long recording is closed on one platform and being withdrawn on the other, leaving the install-free half iPhone-only before the microphone question is even asked.
- **A seat objected:** the entry banks the currently higher net from the United States external-link position as settled, when the same evidence records that position as provisional, since the appeals court has said Apple may charge something and the district court has not set the fee.

G2 , Unmetered capture, priced per finished meeting

- **Proposed by:** the domain generalist, the wide lens on consumer voice productivity software.
- **Who pays:** the individual inside the app store, but for consumable credits against delivered meetings rather than a flat month, at a commission cost of an all-in 15 percent on an auto-renewing subscription against up to 25 percent plus the billing fee on a non-subscription sale.
- **What kills it:** record one 60-minute in-person meeting on a named current handset and run the whole on-device path on it; if the transcript is not finished inside the walk-back window at a battery and thermal cost a person will accept, capture is a cloud bill after all and the unmetered half cannot be funded.
- **Audit verdict:** Yellow, no duplicate and no pivot; the gaps are that no fully local product is documented running on a phone rather than on a laptop, that the only battery figure that circulated has been struck out with nothing put in its place, and that charging for a delivered summary rather than for a recorded hour has no instance anywhere in this project's map of how consumer products here get paid.

Every free tier in this field meters the thing the category is bought for and the thing that is about to cost nothing: transcription dominates marginal cost at 0.21 to 0.62 US dollars per recorded hour while the summarisation call sits at 0.002 to 0.10, two to four orders of magnitude below it, and moving recognition onto the phone takes the metered half to roughly 0.05 an hour or to zero, which Apple's on-device speech framework from iOS 26 and Android's on-device generative stack make available to a third party now. So the packaging the whole shelf shares is inverted: capture and the transcript are never metered, never capped and never lost behind a paywall, and money attaches to the finished summary and action items delivered under the deadline, which is the moment the value lands. That fixes the loss-making shape the funding map works through, where 100 free signups consume up to 200 US dollars of inference a month while the three to five who convert generate 45 to 75, and where the field's other three answers each churn exactly the users you want: minute caps (Otter at 300 a month, Wave at 30, Notta at 120 with a three-minute file cap), history caps (Granola at 30 days) and summary caps (Fathom's first five calls). The counter-evidence is not smoothed away: zero marginal cost is not a moat, it is the mechanism by which the price of the core function goes to nothing, which is the argument for putting the price on the delivery rather than on the recording.

- **A seat objected:** the commission arithmetic here does not reconcile with the two other entries on the same rail, which price the store at the 15 percent small-business rate while this one asserts up to 25 percent plus a billing fee for a consumable, and the whole packaging argument turns on that gap while the second rate carries no source.

T5 , The remorse window

- **Proposed by:** the consumer subscription go-to-market seat.
- **Who pays:** the individual who has just been refunded the device price, on a direct subscription sold away from the app store, where the currently zero-commission United States external-link position applies and is provisional, since the appeals court has said Apple may charge something and the district court has not set the fee.
- **What kills it:** monthly search volume on those exact queries in the target market; if the volume is a rounding error against a shipment base of over two million units the channel cannot carry acquisition, and one keyword pull settles it.
- **Audit verdict:** Yellow, no duplicate and no pivot; the gap is size, because the hardware vendors' return rate, churn rate and repeat-purchase rate are all undisclosed, so the segment could be any size at all.

This is the one segment whose willingness to spend on this job is demonstrated rather than borrowed and it is publicly reachable, because the funding map records retail and marketplace listings as where the category's public buyer evidence lives, in Amazon and Best Buy verified-purchaser reviews, and the hardware vendors run a 30-day no-questions return policy, which means they model remorse and will not disclose the rate. The channel is timed by intent rather than sorted by demographics: acquire on the return, replacement and alternative queries, and on comparison content aimed at somebody

holding a device they have decided against, at the one moment when 159 dollars is back in their pocket and the problem that made them spend it is still not solved. The same review corpus supplies the vocabulary to write with, since the sixty-four stated purchase reasons this project works from were pulled out of it. The discipline that makes the channel affordable rather than a money pit is that 28.1 percent of those statements are about capturing a cellular telephone call, which no application on either platform is permitted to do, so the copy has to disqualify that buyer at the click rather than at the refund.

- **A seat objected:** calling this segment's willingness proven rather than borrowed overstates it, because the largest block of those stated reasons is call capture that no application may do, and the return rate is undisclosed, so the pool is reachable and unsized rather than proven.

T1, Sell readiness, meter the hours

- **Proposed by:** the consumer subscription go-to-market seat.
- **Who pays:** the individual meeting-goer on a direct subscription sold away from the app store, a web checkout the developer bills itself through a merchant of record, with in-app purchase kept only for storefronts where the off-store route is unavailable.
- **What kills it:** compare recorded hours per active user under a metered plan against a flat plan in the same cohort; if putting a visible meter on the price makes people ration what they record, it kills the recording habit the rest of the product depends on, whatever it does to margin.
- **Audit verdict:** Yellow, no duplicate and no pivot; the gaps are that putting the meter on the price list has no instance in this category, so the demand side is unevidenced, and that the free-user arithmetic it answers is itself borrowed from consumer productivity software generally with no footnote from this category.

The category prices flat at 8 to 19 US dollars a month for an individual tier against a cost metered by the recorded hour at 0.21 to 0.62 US dollars assembled, dominated by transcription with the summarisation call two to four orders of magnitude below it, and on the funding map's own worked example a 15 dollar subscription netting about 13 after commission breaks even near 32.5 recorded hours a month, so a user recording two to three hours a working day costs the vendor more than they pay. That inverts the standard consumer playbook, since the engaged user is the loss and the dormant subscriber is the margin, and every growth tactic that drives usage makes the accounts worse. The packaging that matches the cost charges a standing fee for readiness, meaning the recorder is armed, the guarantee is live and the archive is kept, and meters the recorded hour on top of that in packs that do not expire, so the heaviest user becomes the best customer instead of the worst one. The evidence that it could work is that the whole shelf already meters privately and incoherently, with seven vendors each capping a different quantity, 300 minutes a month, 100 a week, 120 with a three-minute file cap, 30 a month, 800 lifetime, five summaries, 30 days of history, and nobody has put the meter on the price list where the buyer can see it and buy more of it, while the hardware buyer already lives under one, since the 159 dollar device buys 300 transcription minutes a month and no more.

- **A seat objected:** this entry and the unmetered-capture entry are opposite prescriptions for the same unit-economics inversion, a visible meter against never metering capture at all, and both are filed as the same product without being marked mutually exclusive, so a mechanical count can rank both high and buy two contradictory price lists.
- **A seat objected:** from the seat that proposed it, that treating a visible meter as compatible with the recording habit and naming habit rationing only as the kill test misses the likelier failure, which is the funnel cost, while the unmetered-capture entry already carries the version of this argument that puts the price on the delivered outcome.

T2, One receipt a year, issued by us

- **Proposed by:** the consumer subscription go-to-market seat.
- **Who pays:** the individual, on a personal card, on the direct off-store rail and reimbursed by an employer afterwards; the store route cannot serve this at all, because the store rather than the

developer is the merchant of record, so the buyer receives a store receipt rather than an invoice carrying their employer's name and tax registration.

- **What kills it:** ask verified purchasers of the dedicated recorder whether they expensed the 159 dollars and whether their employer reimburses software bought on a personal card; if the device was bought personally and never expensed, or bought on a company card through procurement, the wedge is a phantom.
- **Audit verdict:** Yellow, and the question is unusually cheap to put because the buyer is publicly identifiable from retailer reviews; the gap is that the expensing behaviour is asserted from a description of the buyer and no rate, survey or single verbatim buyer statement about expensing exists anywhere in this project's evidence.

The funding map names the hardware rail's real advantage in one line, that the cash arrives before the service does, which is how a company that has shipped over two million units has raised under 6 million US dollars of disclosed funding, and the open question underneath it is what pays for a software entrant's acquisition once the 159 dollars is removed. The only answer available to software is the same one the device uses, which is to take the year in a single transaction at the moment the buyer was already reaching for a card. The second half is the receipt itself: the same map describes the device buyer as an individual who frequently expenses the purchase without going through their organisation's technology function, and records Japan's tax code instantly expensing a recorder priced under 100,000 yen, a demand subsidy that reaches objects and not software, so an annual invoice issued by our own merchant of record and carrying the employer's name and tax registration is the artifact that clears the same expense system a device receipt clears. One correction already on this project's record governs how any of this may be said out loud: annual prepay is a cash-timing and receipt mechanism, never a discount claim, and an annual figure is never set against anybody's month-to-month figure.

An adjacent product for the same buyer

ID	Novelty	Defensibility	Time-to-market	Strategic fit	Total
S5 (folds V3)	3.7	2.3	4.1	4.0	14.1
G4 (folds T4)	3.4	2.7	2.7	3.3	12.1

S5 (folds V3) , Action items that abstain, with the audio behind them

- **Proposed by:** the speech and voice machine learning seat.
- **Who pays:** the individual meeting-goer on a consumer subscription inside the app store.
- **What kills it:** on the sentence-level action-item labels of the AMC-A corpus, or on the wrong-participant labels of the Kirstein annotation set, measure owner-assignment precision as a function of abstention rate; if precision at every rate that still leaves a usable number of items is no better than the non-abstaining baseline, the surface is decoration.
- **Audit verdict:** Yellow, and a declared pivot on this project's three-outputs rule, because a playable evidence span is a fourth deliverable; marked a duplicate of the grounded-attribution group and kept as its strongest statement, with the gap being that the licence status of the two action-item corpora is unknown and the abstention curve cannot be measured without them.

Action-item extraction is the one output in this category with no accuracy number in any product or any paper, and its named failure mode is a fluent, plausible task attributed to somebody who never agreed to it. The failure modes are asymmetric, since a missed commitment leaves the user where they started while an invented one actively misleads them, so precision is weighted far above recall and the product surface that expresses that operating point is abstention: below a confidence threshold the system reports an unassigned commitment rather than naming a person, and every item it does assign carries a timestamped span the user can play. Assigning owners from what was actually said in the room, and abstaining when unsure, avoids voiceprint enrolment entirely, which the technology

survey describes as the safe harbour against named biometric identification under the Illinois Biometric Information Privacy Act at 740 ILCS 14/15(b), and folded in is the counsel's framing that a wrong name is inaccurate personal data about a person who never installed anything, carrying an accuracy duty under GDPR Article 5(1)(d) and a rectification duty under Article 16 owed to the person recorded rather than to the user who recorded them. The supervision exists, which is unusual on this slate: the AMC-A corpus carries 1,506 sentence-level action items over 424 meetings at 0.47 annotator agreement, AIMU 318 actionable turns over 21,035 turns of 22 ICSI meetings, AMI 381 items over 101 meetings, and the Kirstein set 175 machine summaries with human error-type labels at Krippendorff's alpha 0.76 to 0.83, the only material annotated for the exact error this suppresses, while the automatic metrics reward that error, with perplexity at +0.44 on wrong speaker references and BLEU at +0.35 on hallucination (arXiv 2404.11124).

- **A seat objected:** from the seat that proposed it, that the Kirstein set labels wrong-participant attribution as an error type without establishing why it happens, so the claim that the model binds a task to whoever produced most of the words describing it is a hypothesis stated as a finding, and the abstention threshold has to be calibrated against the labels rather than against that story.

G4 (folds T4) , The recap the room receives

- **Proposed by:** the domain generalist, the wide lens on consumer voice productivity software.
- **Who pays:** the same individual on the same subscription; the rail does not change, the acquisition cost does, because every recap places the product in front of everyone who was in the room.
- **What kills it:** put a "who was in the room" step in front of ten real in-person meetings; if users skip it there is no recipient list, and what remains is another notes archive nobody opens twice.
- **Audit verdict:** Yellow, and a declared pivot on the three-outputs rule, since the recap is a fourth deliverable; marked a duplicate of the recap pair, which was folded into it, with the gap being that the payload it sends is unmeasured across the entire field and the recipient list is unsolved.

The bot competitors get the attendee list free from the calendar invitation, which is exactly why Otter and Fireflies can send a recap after a call and an in-person recorder cannot, and the meeting this project is about has no invitation, no attendee list and no shared artifact, which is the structural reason its notes die inside one person's application. So the outgoing recap is addressed to the people who were in the room, assembled from action items that carry owners, and identifying the room becomes the product's real input rather than an afterthought. The supervision for the hard half is thin enough to be worth naming: the AMC-A corpus carries 1,506 sentence-level action items over 424 meetings and 306,846 utterances at 0.47 annotator agreement and the AMI corpus supplies 381 items over 101 meetings, while the technology survey records that no benchmark anywhere scores action items together with their owners, so this output is unmeasured across the whole field and an entrant is not behind on it. Folded in from the go-to-market seat is the observation that the legal cost and the growth loop are the same action, since the announcement is paid for anyway under a binding all-party set of twelve jurisdictions and possibly thirteen, and each delivered recap is a free impression on somebody who attends the exact kind of meeting the product is for, against a fallback acquisition cost the funding map can only quote as borrowed at 20 to 40 US dollars with no footnote behind it; the honest hazard is that sending somebody a recap is proof you recorded them.

- **A seat objected:** calling the recap proof you recorded somebody understates it, because the recipient list is itself personal data about people who never installed anything and every send is a disclosure to a third party carrying a notice duty, not only a consent event.

A different buyer for the same mechanism

ID	Novelty	Defensibility	Time-to-market	Strategic fit	Total
M2	3.9	3.4	4.0	1.9	13.2

ID	Novelty	Defensibility	Time-to-market	Strategic fit	Total
V6	3.1	2.3	4.0	2.7	12.1
S4 + M5	5.0	4.0	1.0	1.9	11.9
P1	3.6	4.0	1.7	1.9	11.2
V4	3.1	2.9	2.4	2.0	10.4
G5	3.0	3.0	2.0	2.0	10.0

M2 , Crash reporting for the microphone

- **Proposed by:** the mobile audio platform engineer.
- **Who pays:** another application developer, on a usage-metered platform licence, the self-serve developer agreement with a negotiated enterprise tier above it that AssemblyAI and Deepgram sell on, paying by volume.
- **What kills it:** open the current event catalogues of the two mainstream mobile crash and performance reporting kits that every application already carries; if audio-session interruption, route change and capture-stall events are already among them, a second software development kit for a signal that arrives free is not bought.
- **Audit verdict:** Yellow, and an undeclared pivot on both the two-taps and the three-outputs rules, since the buyer is an engineering team, the interaction is none, and the deliverable is a structured session record; the evidence gap is that selling a capture-telemetry kit has no instance in this or any adjacent category.

Every product whose value depends on a long microphone session has this problem and none of them can see it, so the layer drops into an application and owns the parts that get people paged: declaring the background audio mode on iOS and the microphone foreground-service type on Android 14 and later, holding the session, subscribing to interruption, route-change and thermal notifications, restarting the service after a memory kill, requesting the battery-optimisation exemption, forcing a universally supported capture rate so that fragmented hardware does not fail silently, and emitting one structured record per session with the audio itself never leaving the device. What the buyer gets that does not exist today is a defect-free completion rate for their own installed base, broken down by handset model and operating-system build, so that a support ticket saying "it stopped recording" becomes a row in a table. The technology survey records that an exhaustive search of developer post-mortems, engineering blogs, mobile-systems literature and public bug trackers found no published capture-survival measurement on either platform. The instrument creates the category baseline as a side effect of being sold, and a public per-handset survival table is the closest thing this project has to the self-liquidating acquisition asset that the 159 dollar object gives the hardware incumbent.

V6 , The room where the announcement is already free

- **Proposed by:** the recording-consent and privacy counsel.
- **Who pays:** the individual professional on a direct subscription through a web checkout away from the app store, the higher-net route on the United States storefront under the current injunction, with the purchase expensed afterwards.
- **What kills it:** check whether the announcement in these roles is already delivered by a scripted preamble read from a form the professional's employer already retains; if the consent record exists on paper in a system the application cannot write into, the in-recording stamp duplicates something and the legal differentiator is worth nothing.
- **Audit verdict:** Yellow, no duplicate and no pivot, and the auditor calls it the cleanest single insight on the slate; the gap is that this project holds no buyer evidence for these roles at all.

Same two-tap capture through the phone's own microphones, aimed at the buyer whose job already requires them to say "I am recording this" out loud before anything else happens: the recruiter running a screening conversation in a room, the reporter taking an on-the-record interview, the adjuster taking

a statement, the field researcher taking informed consent. For every other buyer in this category the announcement is a third interaction and a social cost, which is the objection that makes the compliant shape expensive; for this buyer it is a step they already perform, so the cost is zero, and the artifact it produces, a timestamped in-recording consent stamp bound to the file, is something they currently keep by hand or not at all. The buyer evidence is not overstated: the sixty-four-statement sample was assembled from hardware-owner forums including the sales and consulting communities, so these roles are visible in it, but the stated reasons run the other way, with 17.1 percent naming social discretion and bot avoidance, which is the opposite posture to announcing. What does support the segment is that after the strictest-standard rule every recorder in this market faces the announcement problem, and this is the one buyer for whom it is already solved before the product arrives.

- **A seat objected:** this entry banks the currently higher net from the United States external-link position as settled, when the same evidence records that position as provisional, since the appeals court has said Apple may charge something and the district court has not set the fee.

S4 + M5 , The phone-condition recording programme and the rig that makes the corpus possible

- **Proposed by:** the speech and voice machine learning seat for the programme, the mobile audio platform engineer for the recording rig.
- **Who pays:** unresolved and shared, which is why the audit treats the two as one programme with one payer question: one seat names a United States Small Business Innovation Research Phase I award, recorded as open with no precedent in this category, funding a research question rather than a product and funding a work plan rather than the resulting asset; the other names a speech vendor or research group licensing the collected set, a commercial column nobody has priced because nobody has looked, with public innovation grants as the fallback and the corpus-licensing route itself closed by the voiceprint statutes.
- **What kills it:** one pilot session recording a conversation simultaneously on a phone and on a tabletop microphone array of the kind the existing corpora use, with every participant also on a close-talking reference microphone, all decoded by the same model; if the platform's own resampling drift cannot be aligned by cross-correlation over a single 30-minute session, nothing can be scored against a headset reference, and if the phone stream scores like any other single distant microphone the collection adds a device row to the literature and nobody licenses a row.
- **Audit verdict:** Yellow with no payer named, the tag's exact case, and an undeclared pivot, because a public benchmark corpus is not the transcript, the summary and the action items; the two entries were merged because they share one collection protocol and one unanswered payer question.

The data inventory returns nothing on every claim that needs data, and the structural reason is that no public corpus of any kind was ever recorded on a phone: a recent survey's 36-row table of meeting corpora contains no smartphone entry, and the nearest neighbours are smart glasses, body-worn bin-aural rigs and purpose-built conference devices. The simulation shortcut is closed in both directions, because the field accepts simulated far-field audio for training and rejects it for evaluating a hardware claim, and the NOTSOFAR-1 simulator itself rests on 15,000 physically measured room responses, so the route is a parallel-capture protocol on AliMeeting's topology, an eight-channel array plus a headset on every participant across 118.75 hours in 13 rooms of 8 to 55 square metres at reverberation times of 0.3 to 0.6 seconds, with one identical recognizer decoding every stream because a different recognizer per device fatally confounds the hardware test, significance from the NIST matched-pairs sentence-segment test, and reference transcripts produced by people listening from scratch because the NOTSOFAR-1 creators found that annotators accept plausible but wrong machine guesses. What would make the asset ours is the axis nobody has, placement: face up on a table, face down, shirt pocket, trouser pocket and bag, for which this project's evidence supplies a mechanism for all five and a measurement for none. The second payoff is already sized, with off-the-shelf Whisper at 81.6 percent word error rate on the distant microphones of the LOTUSDIS corpus falling to 49.5 after fine-tuning on distance-diverse overlapping conversational data (arXiv 2509.18722), and synthetic mixtures alone giving 16.0 percent time-constrained speaker-attributed word error rate on AMI's single distant microphone and 20.1 on NOTSOFAR-1 single-channel before any real in-domain audio moved both (arXiv 2605.15442,

measured with ground-truth speaker segmentation supplied), while the three reasons the field avoids phones, automatic gain control breaking the linear amplitude relationship between channels, clock drift with no shared word clock, and firmware tuned for a talker at roughly 30 centimetres to 1 metre, are platform choices with documented workarounds rather than laws of nature.

- **A seat objected:** the kill test is entirely acoustic and omits the instrument that actually gates the corpus, a participant release covering redistribution and commercial licensing, without which the collected audio is unlicensable under the very voiceprint statutes this entry cites, and the metered-inference pivot downstream of it then has nothing to sell.

P1, Pre-installed on the fleet Google does not reach

- **Proposed by:** the consumer platform strategist.
- **Who pays:** a handset manufacturer, licensing the transcription engine for a per-device fee or a revenue share, an architecture the funding map records as having no instance in this category, funded the way the platform owners fund their own, as hardware margin and ecosystem lock-in with an acquisition cost of zero and no store commission.
- **What kills it:** open the current flagship of one of the four named manufacturers and see whether its built-in recorder already returns a transcript, a summary and a distinct action-item list; the claim that Samsung's own recorder does this is carried as uncited and unverified, so one handset settles it, and if the three outputs already ship then the manufacturer has nothing to buy.
- **Audit verdict:** Yellow, and an undeclared pivot on the rule that the hardware buyer is the customer, since the payer is a handset manufacturer; two gaps, that the licensing rail has no instance in this category and the one adjacent case runs the other way, and that the load-bearing technical claim is an inference nothing in this project's evidence establishes.

The free pre-installed competitor is not evenly distributed and the competitive map says so in its own Google profile: Pixel Recorder does the whole job offline with speaker labels since its version 4.2 and is Pixel-only, so it does not reach most Android users, and the developer-facing generative mode of Google's on-device machine learning kit is hardware-gated to Pixel 10 and Pixel 11, which leaves everybody else shipping Android with the demand and none of the feature. The sharper half is that the four manufacturers whose power managers kill a third-party recorder after screen lock, each behind a different multi-step settings path that a major update silently reverts, are the same four parties who could make that recorder un-killable by shipping it inside the system image. A capture path integrated below the application layer is exempt from its host's own process killer, needs no store review and no background audio declaration, and arrives pre-installed rather than through the ranking fight the competitive map calls undifferentiated crowding. The proof-of-concept shape is a capture-survival comparison on one manufacturer's handset, the same recorder run as an installed application and as a privileged system component over long sessions, passing on a survival gap large enough that the manufacturer is buying a defect fix rather than a feature.

- **A seat objected:** the claim that Pixel-only gating leaves everybody else on Android with the demand and none of the feature sits against the competitive map's own list of built-in recorders, which already carries Samsung's as reporting transcription and summarisation, so the largest fleet in the addressable set is the one most likely to have built this already.
- **A seat objected:** it asks four handset makers to accept inside their own system image exactly the wiretap liability the announcement entry argues platform owners structurally refuse, and its kill test never puts that question to them.

V4, The works council package

- **Proposed by:** the recording-consent and privacy counsel.
- **Who pays:** the employer, on team and enterprise seats invoiced direct, which the funding map identifies as the growth engine of every scaled incumbent and the only route in this field with expansion in it.

- **What kills it:** check whether Jamie, which already sells European-only processing and immediate audio deletion into the German market, ships a co-determination template with its deployment; if the incumbent closest to this buyer already hands one over, the package is table stakes rather than a wedge.
- **Audit verdict:** Yellow, with the legal basis verified word for word and the gate real; the gap is that the priced gate on this rail is a security review, so the package removes one barrier and leaves the expensive one standing, and no deployment package has ever been sold or priced separately anywhere in this category.

In Germany the purchase is gated before first use rather than litigated after it, since the consent-law research reports the Federal Data Protection Act section 26 and the Works Constitution Act section 87(1)(6) as requiring a co-determination agreement with the works council before an employer may deploy software of this kind, absent which the deployment is illegal and the council can compel the employer to block it. That blocker sits directly in front of the only compounding revenue route this project has, and it is a document problem rather than a technology problem. The product is the application plus a deployment package built to be signed rather than negotiated: a written retention window, per-participant refusal implemented as a behaviour of the software rather than asserted as a policy sentence, since GDPR Article 6(1)(a) consent must be obtainable from each participant independently and refusible without professional detriment, no model training as a contract term, no biometric enrolment as an architectural fact, and an Article 28 data processing agreement in the box. That the gate is real and closes hard is already recorded in Chapman University's August 2025 blanket ban on Read AI over data security, privacy and consent concerns, and the honest tension is that the individual buyer asks for processing on their own device and not for certifications, which is a different evidence base and a different sale.

G5 , The visit note that lands in the record

- **Proposed by:** the domain generalist, the wide lens on consumer voice productivity software.
- **Who pays:** the sales or field organisation, on seats invoiced direct, away from both app stores and therefore clear of store commission entirely.
- **What kills it:** ask three sales-operations leads whether they would deploy an application that records a customer conversation on a company-managed handset; a blanket legal or device-management refusal ends this before any question of preference is reached.
- **Audit verdict:** Yellow, and a partly undeclared pivot, because it declares the charge against the rule that the hardware buyer is the customer but does not name the fourth deliverable it adds; marked a duplicate of the employer-seat pair and kept as its strongest statement, with the gaps being that no priced instance of this distribution route exists in this project's evidence and that the procurement question is unresolved.

The outputs go into the customer relationship management (CRM) record rather than into a notes archive, and the organisation buys because notes typed in by the representative are the known data-quality hole in every CRM deployment. This is the one buyer where the in-person meeting somebody walks out of has money attached to the write-up, and where no competitor can reach: a bot cannot join a customer visit and a conferencing platform has no stream to hand over, which is why Claap, Zocks and FinMate all built this shape for calls and none of them for the room. The evidence that this is the same person the seed describes is that the buyer sample was drawn partly from the sales and consulting communities, the verbatim call-capture statement comes from a sales representative, and the buyer study describes the purchase path as an individual buying one device outside their organisation's technology function, proving it works, then acquiring units for colleagues. The gate is a security review, SOC 2 Type 2 certification plus a data processing agreement, stated data residency and corporate login, which is a different body of evidence from the one that convinces the person pressing record.

The pivot tier

A pivot here is an idea that contradicts one of this project's own scope rules: the product ships as software on a device the customer already owns, the phone's own microphones are the capture path, the interaction is two taps, the deliverable is a transcript, a summary and action items, the meeting is one the user physically walks out of, the outputs are ready within the few minutes after it ends, and the target is the person who buys or would buy a dedicated recorder. These ideas are kept and argued rather than smoothed away, each one named with the rule it breaks, because an idea that contradicts the thesis is evidence about the thesis.

ID	Novelty	Defensibility	Time-to-market	Strategic fit	Total
M4	3.9	2.9	2.3	2.0	11.1
T6	1.3	2.7	2.1	2.0	8.1
V5	1.1	1.0	5.0	1.0	8.1
P2	2.3	3.4	1.0	1.0	7.7
S6	3.1	2.3	1.0	1.0	7.4
G6 (folds F6)	2.1	1.3	1.7	1.0	6.1

M4 , The meeting is the array

- **Breaks:** the two-taps rule. The person who called the meeting taps start and stop, and every other participant performs their own start on their own phone, which is a third and fourth tap in the room and a coordination cost the original idea does not carry.
- **Proposed by:** the mobile audio platform engineer.
- **Who pays:** the team, on seat-based invoicing, because a product that only works when three people in a room run it is natively a team purchase.
- **What kills it:** record one conversation on two phones at different seats and check whether the per-utterance loudness ratio between them separates the talkers; if each handset's automatic gain control moves independently and fast enough, the ratio carries no seat information and nothing downstream survives.
- **Audit verdict:** Yellow, with the pivot declared with unusual precision, since the seat rules out the techniques independent clocks forbid before anyone can claim them; the gap is the surviving claim, because the energy-ratio attribution cue appears in no evidence this project holds.

The platform will not give one phone four channels, but it will give four people one channel each, and the claim has to be precise about what that buys: independent phones share no hardware clock and drift measurably over a session, which destroys the microsecond phase alignment that beamforming needs, so guided source separation and any classical beamformer are out and must not be promised. What alignment after the fact does support is everything that needs tens of milliseconds rather than microseconds, since cross-correlation alignment is the accepted method for independent devices in the CHiME-6 challenge recipe, and once two streams are aligned at that resolution the per-turn energy ratio between phones is a direct speaker-attribution cue, while channel selection by envelope variance, a CHiME-7 baseline technique, picks whichever phone was nearest the person talking. Every microphone is still a phone's own, so the no-new-hardware and phone-is-the-microphone rules both hold and only the interaction rule breaks. One side effect is worth as much as the acoustics: each participant starting their own capture is consent expressed as an action rather than as a disclosure screen, which speaks to the twelve-jurisdiction all-party set and to the covert-capture theory in *Chamberlain v. Granola*, a case brought against a product that processed on the device, over how the capture was presented rather than where the computation ran.

T6 , Free at the person, paid at the company

- **Breaks:** the rule that the hardware buyer is the customer. It leaves the no-new-hardware, phone-is-the-microphone, three-outputs and in-person-meeting rules intact.
- **Proposed by:** the consumer subscription go-to-market seat.
- **Who pays:** the company, on team and enterprise seats invoiced direct at 20 to 98 US dollars per user per month across the field's published price lists, billed away from both stores; the person who presses record pays nothing, ever.
- **What kills it:** count how many individual installs share an email domain; if the installs arrive one per company there is no expansion path, the free product is pure inference cost, and the cohort reverts to the loss-making one the funding map has already worked through.
- **Audit verdict:** Yellow, pivot declared with the right rule named and both of its stated costs load-bearing; the gap is the free half, since giving the individual product away costs nothing only on a fully on-device path that is unmeasured on any phone, and its kill test cannot run before the free product exists, which makes it the latest-firing kill on the slate.

The funding map's first emerging pattern is that every revenue route that compounds in this field ends at an organisation and not at a person, and the evidence that opens the compounding route is a security review, SOC 2 Type 2, ISO 27001, a data processing agreement and stated data residency, while the sources say the individual and small-team buyer asks for something else entirely, processing on their own device, which is two different products' worth of evidence. Set that beside the crowding pattern, nine or more vendors selling the same three outputs from the same rented models at 9 to 19 dollars, three of them giving the core interaction away and both operating systems shipping it pre-installed and free, and a consumer subscription cannot pay an acquisition cost out of what one person spends in a category whose own price floor is zero. The pivot reorders the business: the individual product becomes the acquisition channel, given away and built so that giving it away costs nothing, since on-device transcription sits near 0.05 US dollars a recorded hour and a fully local pipeline at zero, and the revenue is the seat contract sold into the companies where those free individuals turn out to cluster, which is the motion this field's winners run, with Otter growing on Business and Enterprise, Read AI valued on Fortune 500 penetration rather than on the 100,000 consumer accounts it adds weekly, and Granola's Series C raised on an enterprise-context thesis. Two honest costs: the seed's customer stops being the payer and becomes the distribution, and the certification toll is paid in full before the first seat contract is signed, in a currency the free user does not value and will not fund.

V5 , The record you speak yourself

- **Breaks:** the three-outputs rule. The first output stops being a transcript of the meeting and becomes a transcript of the user talking about the meeting, and this project's charter excludes a dictation product by name.
- **Proposed by:** the recording-consent and privacy counsel.
- **Who pays:** the individual, on a one-time lifetime purchase, which the funding map opens only where marginal cost is near zero, and a single close-talking speaker is the case where inference on the phone is genuinely adequate.
- **What kills it:** over a run of the founder's own meetings, count how many actually get dictated within the walk back; if most are skipped the premise fails outright, because the product asks the user to remember the meeting, which is the job they were trying to hand off.
- **Audit verdict:** Yellow, pivot declared, and the only idea in the round with no non-user plaintiff class, no stranger's voiceprint and no consent disclosure, which is a structural property rather than a posture; the gap is the one the seat states against itself and it is fatal-sized.

Do not record the room at all: on the walk back the user taps start, says what happened, taps stop, and gets back a transcript of their own words, a summary of it and the action items in it. Every legal exposure in this project's evidence is defined by capturing the voice of somebody who never agreed to anything, with four pending United States actions on record (In re Otter.AI Privacy Litigation arising from Brewer v. Otter.ai Inc., Cruz v. Fireflies.AI Corp., Basich v. Microsoft Corp. and the Chamberlain action against Granola), the German Criminal Code section 201 making it a criminal offence for a participant

to record the non-publicly spoken word at up to three years, French Penal Code Article 226-1 requiring all-party consent with a proportionality test from the French data protection regulator on top, and every one of the twelve in-person all-party jurisdictions regulating the capture of another person. It is also the acoustically trivial case, one close-talking speaker rather than the far-field multi-talker condition the whole far-field effort exists to survive, which is why it runs on the handset at zero marginal cost and finishes inside the walk back. Argued honestly it gives up a great deal: it serves none of the 64 stated hardware purchase reasons, it cannot touch the 28.1 percent cellular-call segment any more than the original can, and it loses the verbatim record that is the thing a meeting recorder is bought as insurance for.

P2 , Call recap licensed into the carrier's dialer

- **Breaks:** the in-person-meeting rule, because the audio is a telephone call rather than a room somebody walks out of, and the phone-is-the-microphone rule, because the capture path is the call stream rather than the handset's microphones.
- **Proposed by:** the consumer platform strategist.
- **Who pays:** a mobile network operator, bundling it into the native dialer as a churn and retention feature funded out of subscription revenue, a route recorded as live and reachable only by the carrier, bypassing both stores and every commission schedule.
- **What kills it:** check whether any Western carrier ships call recording in its own dialer today, and note the second cheapest observation, that the iPhone dialer belongs to Apple, so the route reaches Android subscribers only and a carrier will not market a feature it can offer to part of its base.
- **Audit verdict:** Yellow, both broken rules named, and it uses the call-capture finding the right way round by going after that segment from the only layer where it is servable; the gap is that no third-party application has ever been carrier-bundled in this category and the one working instance is a carrier serving itself rather than a supplier being bought.

The funding map closes third-party call capture as structurally dead on both platforms, iOS exposing no public interface to the audio of an active call and Google Play banning the accessibility-service route since May 2022, and it names the one entity for whom it is not closed, whoever owns the dialer. The competitive map's SK Telecom profile is the working instance, an assistant built into the native dialer that transcribes, translates and summarises for over ten million subscribers with no download and no app store, described there as the shape of the threat no application can answer, because it does not compete for the download. Against that sits the buyer evidence: 18 of 64 explicit purchase-rationale statements, 28.1 percent and the largest single category, are about capturing a cellular telephone call, quoted verbatim as "Otter doesn't record phone calls (thanks, iOS)", and those buyers did not walk past an equivalent application, because there is none and can be none. The proof-of-concept shape is a single carrier's answer to whether it can originate call recording for its own subscribers in its own market, where pass is a yes with a named legal basis and kill is anything softer.

S6 , Phone-condition recognition sold by the hour

- **Breaks:** the rule that the hardware buyer is the customer, since the buyer is an engineering team and never a meeting-goer, and with it the two-taps and three-outputs rules, because a metered programming interface has no interaction and produces neither a summary nor action items.
- **Proposed by:** the speech and voice machine learning seat.
- **Who pays:** another developer or an enterprise data team, on usage-metered speech licensing, self-serve with a negotiated tier above it, the way AssemblyAI and Deepgram sell it.
- **What kills it:** decode the same phone-captured meeting recordings through two published general interfaces and through the tuned model and compare the error rates; if the tuned advantage on the phone condition is smaller than the price gap to a general vendor, the buyer takes the general vendor.
- **Audit verdict:** Yellow, all three broken rules named, and the payer rail is live with named precedent, which is more than most pivots here have; the gap is that its wedge is the phone condition and the

phone condition has no corpus, so it sits downstream of the recording programme and inherits that programme's unresolved state entirely, and its own kill test cannot be run until the recordings exist.

This rail competes on price per minute against companies whose prices are already published, Deepgram Nova-3 at about 0.258 US dollars per audio hour, AssemblyAI at a 0.15 base and recorded at 0.37 with speaker separation and entity detection switched on, Groq quoted near 0.04, so a general recognizer sold this way is a commodity and we would lose on price. The only defensible wedge is the condition rather than the model: none of those vendors has ever been measured on a phone in a room, their published accuracy figures name neither a corpus nor a microphone condition, and a supplier that can state a time-constrained speaker-attributed word error rate on phone-captured far-field meeting audio at a named placement against a named reference sells the one thing nobody else in the supply chain can state. The buyers who would care are the ones who cannot reach the condition themselves, the wearable and card-recorder vendors whose devices sit in the same acoustic position and the application developers whose products already record through the phone. Selling inference is the version that survives, because licensing the recorded audio itself is closed by the voiceprint statutes and by the category's own contractual promises never to train on customer data.

G6 (folds F6) , The accessory we do not build

- **Breaks:** the no-new-hardware rule, and with it the phone-is-the-microphone rule, because the audio arrives from an object the customer buys separately.
- **Proposed by:** the domain generalist, the wide lens on consumer voice productivity software.
- **Who pays:** the user pays an accessory maker at retail for the object and pays us for the software inside the app store; the alternative is a per-device software licence paid by the accessory maker, which the funding map lists as an architecture with no instance in this category.
- **What kills it:** find out whether any accessory maker sells a recording button with a documented, open synchronisation interface; Anker's ten-gram bean was built for ByteDance's Lark alone, and if that exclusivity is the norm then a software entrant has no supplier and would have to stock and take returns on the object itself.
- **Audit verdict:** Yellow, both broken rules named, and the seat states the counter-evidence against itself, which is the right handling; the gap is the supplier, since the only located instance was a captive build for a single software owner and the accessory maker's payment route has no instance in this category.

The buyer evidence is the whole argument: all 64 of 64 stated purchase reasons sit in categories the buyer study classifies as either structurally impossible for a third-party application or inherent to owning a separate object, call capture at 28.1 percent, battery preservation at 14.0, physical start-stop tactility at 7.8, and not one buyer mentions audio quality. The field has already run this experiment and it ran the other way round from the hardware incumbents: ByteDance kept Lark as the product and let Anker build the recording bean as an accessory that syncs into the software, hardware treated as an accessory to software rather than as the product, so the software keeps its economics and its rails, the object is somebody else's inventory, warranty and returns, and the customer gets the separate battery and the physical button that 21.8 percent of stated reasons are about. The counter-evidence is stated plainly rather than argued around: this is the second-object pattern that killed Humane and left Limitless owners holding a brick when Meta switched the service off, the competitive map rates that pattern as the one thing this project is currently immune to, and the pivot contradicts the seed's entire ten-times-better claim. Folded in from the founder seat is the passive variant of the object, a start and stop button and nothing else, no microphone, no data connection and no audio path, which turns entirely on whether any documented interface lets a paired accessory start third-party microphone capture on a locked phone.

- **A seat objected:** the passive button is barred before any accessory maker is even asked, because Android 12 and later refuse a background start of a microphone service outright and a backgrounded iOS application cannot activate a recording session, the same refusal class as the error code the capture-survival entries are built around.

What was dropped, and why

Forty-one ideas were proposed and twenty-seven entries were scored. Nothing was erased; every idea that is not a scored entry above was either folded into one that is, or dropped for a stated reason, and both are listed here.

- **Never Lost** (founder seat), the capture guarantee stated first and plainest: dropped, because it adds nothing the two kept capture-survival statements already carry, and its purchase evidence is a block of a buyer sample drawn from people who already own the hardware.
- **The capture receipt** (domain generalist), the same guarantee with the sharpest test attached: dropped for the same reason, but its interruption protocol survives as the cheapest observation on the whole slate, named at the end of this document.
- **Verified capture** (mobile audio platform engineer): folded into the watchdog entry as the pre-flight probe, the one element in that cluster no other seat proposed.
- **A published maximum, with a refund behind it** (go-to-market seat): folded into the published-deadline entry as the move that puts the number into the terms of sale with a remedy behind it, corrected to the off-store rail; its capture-survival half went to the capture-survival group.
- **A Number On The Wait** (founder seat): dropped as the published-deadline entry without its cost, because it does not carry the on-device context limit that makes the post-tap pass hard.
- **Everyone In The Room Can See It** (founder seat): the one red, dropped. It has the phone speak the recording notice and then claims the two-tap interaction survives intact because the announcement is fired by the start tap, when this project's own evidence states in two places that the compliant interaction is at minimum start, announce, stop, and that the cost is the spoken act in the room rather than the finger on the glass. Three other seats propose the same mechanism while paying that cost honestly.
- **The announced recorder** (domain generalist): dropped as the announcement idea without a mechanism the product can own.
- **The consent record the platform will not ship** (consumer platform strategist): folded into the announcement entry for its argument that a platform owner structurally cannot occupy this position.
- **Receipts** (founder seat): dropped, because the playable evidence span it is built on is present in both surviving grounded-attribution statements.
- **No name on a sentence without the sentence** (recording-consent counsel): folded into the abstaining action-items entry for its framing of a wrong name as inaccurate personal data about a person who never installed anything.
- **The Rep Who Walks Out** (founder seat): dropped in favour of the visit-note entry, which names the destination that makes willingness to pay concrete and names the three competitors who built the call version and not the room version.
- **The Button That Is Not A Recorder** (founder seat): folded into the accessory pivot as one line, because that entry reaches the same buyer motive without putting inventory, warranty and returns on our own balance sheet.
- **The recap is the invite** (go-to-market seat): folded into the recap entry, which files the same delivery as a fourth output and therefore as an adjacent product rather than as the same one.
- **The rig that makes a phone-recorded meeting corpus possible** (mobile audio platform engineer): merged with the phone-condition recording programme into a single entry, because the two share one collection protocol and one unanswered payer question, and neither can be judged without the other.
- The two ambisonic entries were deliberately not folded: they propose different products, a separation front end and speaker labels, off one and the same measurement, so the recommendation is to run one recording and keep both downstream uses.

What this slate does not answer

A reader of the pitch arrives holding one line, that we give the person who was about to buy a recording device the recording device they already own, and leaves this slate without a single idea that tests

whether the phone can do the job. The open question of what word error rate a real phone achieves in a real meeting at five named placements, face up on a table, face down, shirt pocket, trouser pocket and bag, is measured by nothing here; the two entries that propose building the instrument to measure it are the two the audit tagged as having no payer. Nor does the slate engage this project's own first kill criterion, which the evidence says fires on its literal wording: the product already exists free, Voicenotes does the exact flow at zero, Fathom and Granola and Hedy give the interaction away, and both operating systems ship it pre-installed. Almost every idea here is a reason to prefer our recorder over another application. Two are reasons to prefer an object, and they answer the acquisition question by putting the 159 dollars back. None is a reason for the person holding the device to put it down, which is the sentence the pitch is made of.

The cheapest observation that would most change this ranking is the interruption protocol from one of the folded capture-survival entries: take two shipping recorder applications, one iPhone and one Samsung handset, force four conditions on each, an incoming call mid-session, a screen locked for an hour, the manufacturer's own battery daemon, and a second application requesting the microphone, and count how many sessions come back with complete audio. It needs no corpus, no annotation and no counsel, and it is the only test on the slate whose outcome moves six ideas at once in either direction. If the incumbents survive it cleanly, the highest-rated cluster in the round has nothing to sell against and the slate collapses to the announcement group and the employer-seat group. If they fail it, that cluster becomes the only mechanism in this project that is simultaneously buildable, unclaimed and evidenced, and every other idea should be re-ranked beneath it. Two other observations are comparably cheap but move less: reading Apple's and Google's published policy for their zero-install surfaces settles the record-before-install entry in one paragraph and touches one idea, and opening one named flagship handset to see whether its built-in recorder already ships all three outputs settles the pre-installed entry and the unverified claim about Samsung's own recorder, and touches two.