Recapp

Brief review

This file records what a panel of eight readers said about the project's two-page investor memo.

2026-09-02

COPUNDIT

Prepared by Copundit

AI for Advanced Ideation™

This file records what a panel of eight readers said about the project's two-page investor memo. Each reader was given the same memo and the same five questions to score it on, and each wrote up their objections without seeing what any of the others wrote, so agreement between two seats here means two people reached the same conclusion separately rather than one persuading the other. The panel did not touch the memo. Revising it is a separate step that someone still has to decide to take, and this file is the input to that decision. Where two readers reached opposite conclusions on the same sentence, both are printed as they were written; nothing is averaged away.

Investor verdict: Not fundable as it stands, and the one thing that decides it is the two-handset test on shipping recorder apps: if they lose meetings often enough to matter there is a company here, and if they do not there is not.

The scoreboard

Seat	Mechanism	Evidence	Buyer economics	Timing	Defensibility
Investor	2	4	2	3	2
Founder	3	3	2	3	2
Domain generalist	2	2	2	3	2
Speech and voice machine learning	2	3	2	2	2
Mobile audio platform engineer	3	4	3	3	2
Consumer platform strategist	3	3	2	3	2
Consumer subscription go-to-market	3	4	2	3	2
Recording-consent and privacy counsel	3	3	3	4	2
Mean of eight seats	**2.62**	**3.25**	**2.25**	**3.0**	**2.0**

Defensibility is unanimous at 2, the only lens on which eight readers who could not see each other landed on the same number. Evidence is where they split hardest, from 2 to 4: three seats called the sourcing the memo's best feature because it publishes its own dead claims, while the domain seat gave it a 2 for containing no word error rate, no test set, no diarization error rate and no latency figure of its own. Buyer economics is the lowest mean and no seat scored it above 3. The unanimous 2 is not eight readers being harsh; it is eight readers agreeing with the memo's own sentence that there is no technical moat here, and then finding that the substitute it offers is copyable too.

Must fix

- **The "no per-recorded-hour cost" claim does not survive the on-device session limit.** Lands in the opening, both "Why now" bullets, and "What we would build". Raised by the investor, the founder, the domain generalist, the speech seat, the mobile audio engineer and the platform strategist. The correction already registered against this claim in the evidence base puts the on-device model's session ceiling at 4,096 tokens covering the prompt, the transcript and the output together, against 11,000 to 13,000 tokens for an hour of speech, and the strategist adds that Google's small on-device model caps the same way, which is itself recorded as unanswered. So the phone does not summarise a meeting in one pass and the memo's only margin advantage is unsupported as written. Four of the six ask for the identical repair in nearly the same words: say chunked, or say cloud.
- **The memo prices a cloud pipeline one paragraph after saying it does not build one.** Lands in "Why now" against "The buyer already pays a subscription". Raised by the investor, the domain generalist, the speech seat and the go-to-market seat. The investor puts it as a fork: either the metered hour is the incumbent's problem and an argument for us, or it is ours and the on-device claim is soft, and it cannot be both. The speech seat says the evidence base puts the on-device path near 0.05 dollars a recorded hour with breakeven past 250 hours, and that the path exists on one platform only because Google's is alpha. Go-to-market wants it split by platform and one of the two numbers dropped.
- **The 32.5-hour breakeven does not reproduce from the figures printed beside it.** Lands in "The buyer already pays a subscription". Raised as must-fix by the founder, the speech seat and the platform strategist, and as high priority by the domain generalist. Founder and speech both get 24 to 71 hours from 15 dollars over the printed 0.21 to 0.62 range, and both find that the source's 32.5 quietly assumes about 13 dollars net of store commission and a 0.40 dollar median hour. The strategist gets 37.5 from 15 over 0.40 and says the commission line is missing and its rate is live, with the United States external-link fee unset after the May 2025 injunction, the European schedule moving on 1 October 2026 and Google Play's terms unestablished and recorded as a gap. The generalist gets 36 from the midpoint.
- **There is no acquisition channel anywhere, and removing the hardware removes the one the category has.** Lands in "The buyer already pays a subscription", and in no other paragraph because there is none. Raised as must-fix by the go-to-market seat and the investor, and as high priority by the founder and the domain generalist. The investor's version: the 159 dollars is what pays to find the customer, which is why a vendor shipping over two million units raised under 6 million, so the paragraph reads cost-neutral to the seller and it is not. Go-to-market's is different: the object was the channel, cash-positive at the point of sale, and every scaled incumbent in the category grows on team seats, which the charter forbids, so the memo must say how the first paying users are found. The generalist notes this is the category where acquisition cost is what kills companies.
- **The single-channel sentence is wrong on the platform facts and borrows its number from a different experiment.** Lands in "Here is what a search does not return", the page-one close. Raised by the speech seat, the mobile audio engineer and the founder, for three different reasons, all kept. Mobile: iOS has released two channels since iOS 14 and iOS 26 adds directional first-order ambisonics, the evidence base forbids the word "permanently", the Android half is untested per manufacturer, and "the recorder does not" is false against a competitor device that is itself mono; the only true sentence is that no raw per-microphone channels are released. Speech: the measured 11.4-point channel penalty is one microphone against a seven-microphone tabletop array, not a phone against a recorder, which is a correction this project has already registered. Founder: one of the three things that would unstick the channel question shipped in the very release the memo leads on, so this is dead pending a test, not dead.
- **The Android screen-lock kill is stated flat where the source it came from says the opposite.** Lands in "Here is what a search does not return". Raised by the founder, the domain generalist and the mobile audio engineer. "Four handset makers kill a correctly declared recording service at screen lock" outruns its source, whose own state matrix records screen lock as continuing, and whether a locked foreground session counts as backgrounded is an open question the business sits on. The generalist adds a second defect in the same paragraph: the iOS call-suspension finding

is a research response, not “Apple’s own documentation” as the memo attributes it. Mobile adds a third: the named Android maker’s kill is its sleeping-apps tier and memory eviction, not screen lock at all.

- **The two named tests cannot produce the number the memo stakes itself on.** Lands in “No amount is set”. Raised as must-fix by the mobile audio engineer, and as high priority by the founder and the domain generalist. Mobile: 500 sessions cannot resolve a 99.0 percent bar to a point when failures cluster by handset model and manufacturer skin, and the stimulus chosen for the one Android handset is the wrong one for that handset’s actual kill mechanism, against a four-manufacturer claim. The generalist: 500 sessions cannot separate 99.0 percent from about 97, and the unit that varies is the handset, not the person, so recruiting 50 people is measuring the wrong axis. The founder: one iPhone and one Samsung cannot yield a number across handsets, operating-system versions and power settings, which is exactly what the defensibility claim requires.
- **“Over two million shipped” is unattributed and conflicts with the only figure anyone fetched.** Lands in the opening. Raised as must-fix by the founder and the domain generalist, and as high priority by the platform strategist. The single fetched source says more than 1.5 million, reported January 2026, and the competitive map records that figure over a device line rather than over the card the memo describes. The generalist: name the tracker or use the fetched number. The founder adds that the memo never says whose two million these are, and that the incumbent goes unnamed in its own opening sentence.
- **Speaker attribution is never mentioned, and one of the three promised outputs depends on it.** Lands in the opening and “What we would build”. Raised as must-fix by the domain generalist, and as high priority by the speech seat and the privacy counsel, each for a different reason. Generalist: every named rival ships speaker labels, neither platform exposes speaker separation to third parties, and the memo does not say whether the transcript names speakers or where that computation runs. Speech: action items need owners, the registered a registered correction records both platforms withholding speaker labels from third-party apps, and on-phone speaker separation has no measured real-time factor anywhere, which is recorded as an open question. Privacy: speaker separation is precisely what two of the pending complaints plead as voiceprint collection under the Illinois biometric statute, and the scope decision is recorded as a decision this project has not made.
- **The memo retires the hardware buyer and never names the buyer that replaces them.** Lands in “So we asked why” and “The buyer already pays a subscription”. Raised as must-fix by the go-to-market seat, and as high priority by the investor and the privacy counsel. Go-to-market: the three stated reasons quoted, 28.1 plus 20.3 plus 17.1 percent, are unserviceable or worse in software, so at most a third of that demand is reachable, and the memo then prices a generic subscriber it never identifies. Investor: the paragraph reports three of six categories, while the source’s own capability matrix puts all 64 statements in rows an app cannot supply, and a partner finds that. Privacy: the call block is also the most regulated medium, since some states are all-party by phone and one-party in person.
- **The capture-tightening bullet is false as stated, and its evidence stops before the release the memo leads on.** Lands in “Why now”, second bullet. Raised as must-fix by the platform strategist and the mobile audio engineer, and marked down in the timing score by the investor, the generalist, the speech seat and the go-to-market seat as well. Strategist: “every documented background-audio change tightened capture” has a counterexample, since Android 15 exempted microphone foreground services from the six-hour cumulative timeout, so the claim has to narrow to background starts. Mobile: the cited window ends at Android 15 and iOS 17 while the product requires iOS 26, so audit the newer releases or say plainly that the trend line stops where it stops. The investor calls it a five-year drift and the generalist says the bullet was equally true in 2023, which is the half the memo calls the business.
- **A suspended or killed app cannot deliver the in-room warning the memo sells as its differentiator.** Lands in “What we would build”. Raised as must-fix by the speech seat and the mobile audio engineer, and cited in the mechanism score by the investor, the generalist and the go-to-market seat. Mobile: a process killed by a manufacturer’s sleeping-apps tier, a freeze policy or a memory reclaim runs no code and warns nobody, and nothing restarts itself either, so recovery is a

third tap that the memo has already spent on the consent announcement. Both seats name the one mechanism that survives the failure it is meant to report: a pre-armed notification that each healthy heartbeat cancels. Speech: name it or drop the claim.

- **“Start with what is not true” concedes the category and never answers it.** Lands in “Start with what is not true”. Raised by the go-to-market seat and the founder, for opposite reasons, both kept. Go-to-market: the memo explains why we beat the hardware and never once says why anyone pays us instead of the five free products it just listed. Founder: the concession is factually too generous, because neither pre-installed product is established to ship action items as a first-class output and one of them runs on a single manufacturer’s handsets, so all three outputs were given away for nothing.
- **The convergence answer is contradicted by surfaces the platforms have already shipped.** Lands in “There is no technical moat in a wrapper over rented models”. Raised as must-fix by the platform strategist and as high priority by the investor, on two different grounds. Strategist: a unified task manager over active foreground services and per-app battery-restriction toggles are the platform shipping the warning the memo says a platform owner will not ship. Investor: the sentence answers the wrong threat, because the threat is the platform fixing the kill, which deletes the product rather than competing with it. Two further seats note in passing that the microphone indicator and the microphone foreground-service type shipped years ago.
- **Recognition during the meeting is charged to the battery the memo’s premise rests on.** Lands in “What we would build”. Raised as must-fix by the mobile audio engineer and as high priority by the speech seat. Mobile: recording alone is nearly free, but continuous on-device inference costs 25 to 35 percent of a current flagship iPhone’s battery an hour and invites the thermal throttling the evidence floor lists as a buffer-dropping cause. Speech splits the same figure: 10 to 15 percent an hour for recognition alone, 25 to 35 for a full offline pipeline. The memo’s opening argues that the phone already has the battery.
- **There is no store-review gate in the plan, and it comes before both tests.** Lands in “What we would build”. Raised as must-fix by the platform strategist and as high priority by the mobile audio engineer. App Store review guideline 2.5.4 together with the background audio mode is a documented approval risk for exactly this shape, a phone face down with no visible feature that needs the declaration, which is recorded as unanswered. The strategist notes it precedes both of the tests the memo hangs its decision on, so it is the first gate and it is unlisted. The founder seat reaches the opposite conclusion on the same guideline; see the disagreements.
- **“Tap start and put the phone face down” is not what an Android user actually does first.** Lands in “What we would build”. Raised as must-fix by the domain generalist and as high priority by the go-to-market seat. The remedy on record is a per-brand settings walk before the first recording, reverted by system updates, which is recorded as the only remedy anybody has published. Go-to-market reads the same fact as the opening: two taps is already table stakes for a returning user on two of the named rivals, and the unclaimed position is the first run, which is a multi-step gauntlet across the whole shelf.
- **The latency promise carries none of the identifiers the field requires.** Lands in “What we would build”. Raised by the domain generalist. The rival’s turnaround is quoted with its recording length attached, and ours is a phrase about a walk. State the recording length, the measured interval from tap-stop to all three outputs, and whether it runs on the phone or in the cloud, which is one of this project’s own open questions.
- **The 99.0 percent bar counts completion only, so a useless recording passes it.** Lands in “No amount is set”. Raised by the speech seat. A session that runs end to end at 30 percent speaker-attributed word error rate clears the bar exactly as written. Add a quality floor; the nearest published anchor is the single-microphone winner of a recent far-field meeting challenge at 22.2 percent.
- **The consent clause is written as the exception when it is the default posture.** Lands in “What we would build”. Raised by the privacy counsel. “Where every party must consent” reads as a minority condition and the evidence base says the reverse: under Kearney v. Salomon Smith Barney (California, 2006) a meeting with one out-of-state participant inherits the stricter law, so all-party is the national default posture. Do not repair it with a state count, because the eleven-state framing is already a registered correction.

- **“Put the phone face down” markets covert capture, and the category is being sued over that exact pattern.** Lands in “What we would build”. Raised by the privacy counsel. Chamberlain v. Granola, filed in a Northern District of California court in July 2026, sued a locally-processing note-taker over this fact pattern. On-device processing is a go-to-market asset, not a legal shield.
- **The pending actions against this category are absent from the paragraph that surveys it.** Lands in “Start with what is not true”. Raised by the privacy counsel. The paragraph names several of the defendants as free competitors and none of the four complaints against them, one of which is past a motion-to-dismiss hearing. By the memo’s own stated standard, a partner finds this in ten minutes.
- **The survival dataset as described is a commercial collection of non-consenting third-party voices.** Lands in “No amount is set”. Raised by the privacy counsel. Five hundred sessions from fifty people captures people who installed nothing and agreed to nothing. The evidence base gives the lawful shape instead: timestamps, buffer counts, interruption events and watchdog firings, with audio never leaving the device, and it should not be called research, because the same section says that framing invalidates the lawful basis.

High priority

- **The shut-down comparable had two causes and only one is printed.** Lands in “There is no technical moat in a wrapper over rented models”. Raised by the founder and the platform strategist. Vowel died of free native summaries and then a withdrawn term sheet, the competitive map says the order matters, and printing one cause overstates the convergence threat.
- **Digital silence is more detectable than the memo concedes.** Lands in “What we would build”. Raised by the mobile audio engineer, with the speech seat saying the same thing in its notes. Android exposes a client-silenced query and iOS 17 an input-muted property, and a live microphone has a noise floor, so a run of bit-exact zeros is itself a detection; the residual risk is a manufacturer’s own audio layer. The speech seat says one handset logging per-buffer peak amplitude and zero-run length closes the open question.
- **No market is built from the bottom anywhere in the memo.** Lands in the opening. Raised by the investor. Two million cumulative units for one vendor is the only count in the document, with no payer count, no price and no conversion rate.
- **Google’s on-device path is not wholly gated to two handsets.** Lands in “Why now”, first bullet. Raised by the platform strategist. The basic mode runs on Android 12 and above; only the advanced mode is Pixel 10 and 11.
- **“No amount is set” reads as the raise, one paragraph after a price paragraph.** Lands in “No amount is set”. Raised by the domain generalist.
- **No launch geography and no retention posture.** Lands nowhere, which is the objection. Raised by the privacy counsel. German penal code section 201 criminalises the act for a participant, French penal code article 226-1 requires all-party consent, and the European right to erasure requires deleting one named person’s contributions from a recording of a group.
- **The first test omits two stimuli that cause the failure it is looking for.** Lands in “No amount is set”. Raised by the mobile audio engineer: a Bluetooth route change and the battery-optimisation exemption.
- **The wedge recruits for the incumbent.** Lands in “There is no technical moat in a wrapper over rented models” and the closing line. Raised by the go-to-market seat. “Your phone may stop recording” is the hardware vendor’s own closing argument, and survival is a retention asset rather than an acquisition one; the acquisition claim actually available is latency with a number attached, which is where a quoted user left hardware for software, and the memo leaves it unclaimed.

Where the seats disagreed

- **Is a silent audio callback detectable from inside the app?** The memo concedes it may not be. Three seats took the concession at face value and priced it: the investor scored mechanism 2 because detect-and-warn is the product, the domain generalist scored mechanism 2 calling it the exact failure the memo sells against, and go-to-market called it one of the two failures the memo sells against. The speech seat and the mobile audio engineer say the opposite on the same sentence: a live microphone has a noise floor so an exact-zero run is itself a detection, and both platforms expose an explicit mute or silenced signal, leaving only a manufacturer's audio layer as genuinely undetectable. Two seats think the memo's honesty about this is fatal; two think the memo conceded something it did not have to.

- **Can the survival dataset be collected at all?** The privacy counsel's must-fix says 500 sessions from 50 people is a commercial corpus of non-consenting third-party voices, and that the lawful shape is metadata only with the audio never leaving the device. The speech seat's notes say the opposite: those 500 sessions would be the only phone-captured in-person meeting audio in existence, the data inventory records no such corpus anywhere, and transcribing a subset buys an evaluation set that outlives any single operating-system release. One seat wants the audio kept and labelled; the other wants it never to leave the handset. This is the sharpest collision on the panel, because the two proposals cannot both be executed.
- **Is capture survival a moat, and did anyone actually say so?** Two seats pre-emptively rebutted a claim nobody made. The investor wrote that another seat would call capture survival a moat and that it is a publication, compounding only if it drives per-handset capture behaviour. Go-to-market wrote that other seats would call the survival dataset a moat and that it is re-earned every release. Neither prediction landed: no seat called it a moat, the speech seat said it lasts one operating-system release, the founder said it is the right asset but the proposed test cannot produce it, and defensibility came back unanimous at 2. Two independent readers each expected to be the lone sceptic and all eight were sceptics.
- **What is wrong with the convergence sentence?** The platform strategist says the premise is factually false and names shipped counterexamples. The mobile audio engineer and the domain generalist agree it shipped, and mobile adds that the platform owner's real move is to exempt its own recorder, as the two pre-installed products already are. The investor does not dispute the premise at all and says the sentence answers the wrong threat, because the danger is the platform fixing the kill rather than shipping a warning about it. Correcting the fact and answering the right threat are different repairs and the memo cannot make both edits to one sentence.
- **Is the consent announcement a per-jurisdiction mode or the default?** Both seats predicted each other, and both predictions landed. The founder's notes say a privacy seat will argue the announcement kills the two taps, and that twelve jurisdictions, possibly thirteen, is a per-jurisdiction mode rather than a surrender. The privacy counsel's notes say a product seat will read start-announce-stop as a small interface tax, when it is a spoken act in a room of colleagues and it is the default rather than the exception, and its must-fix explicitly forbids repairing the memo with a jurisdiction count because that framing is already a registered correction. The founder is reasoning from exactly the count the counsel says is dead.
- **Does the store-review guideline sink the face-down design or pass it?** The platform strategist and the mobile audio engineer both flag guideline 2.5.4 as a documented approval risk for a background-audio declaration with no visible feature, which is what a phone face down is. The founder's notes cite the same guideline and reach the opposite conclusion: the in-room warning is the visible feature the guideline demands, so the warning is what gets the declaration approved. Same rule, same product, opposite verdicts.
- **Is the free-competitor paragraph too generous or not generous enough?** Three positions on one paragraph. The founder and the platform strategist say it gives away too much, because no first-party recorder is established to ship action items as a distinct output and one of the two named pre-installed products runs on a single manufacturer's handsets, and the strategist calls that the memo's best card. The domain generalist says it gives away too little, because calling two of the rivals free giveaways undersells products that ship an in-person iPhone path and publish a turnaround under three minutes, and they are the comparison rather than a free adjacency. Go-to-market says the concession itself is fine and the failure is that the memo never answers it.
- **Is the single-channel claim dead or open?** The speech seat says the memo should say less and keep it closed: the comparison is unevidenced and the only defensible sentence is that raw per-microphone streams are withheld. The founder says the memo said too much and should reopen

it: three named things would unstick the channel question, one of them shipped in the release the memo cites, so it is dead pending a test and not dead. The mobile audio engineer supplies facts that support both, since the platform releases some channels but not the raw per-microphone ones.

- **Is this product's marginal cost near zero or metered by the hour?** The speech seat says the on-device path is near 0.05 dollars a recorded hour with breakeven past 250 hours, so the printed 0.21 to 0.62 is somebody else's cost structure. Go-to-market goes further and calls zero marginal cost on one platform a pricing weapon rather than a cost line, unlimited minutes against a shelf metering at 300. The investor says the on-device claim is the one that is soft, given the session limit, so the metered hour may well be ours. Three readers looked at the same two numbers and reached three different conclusions about which one describes this company.
- **One prediction that did land.** The mobile audio engineer wrote that another seat would read the phone-call block as permanent, and that on iOS it is an absent interface while on Android it is a store policy dated May 2022, with the durable half being the vibration transducer question. Go-to-market and the privacy counsel both treated the block as settled, one to size the reachable demand down to a third and the other to note the medium is the most regulated one. The prediction was correct and the correction is unabsorbed.

The seats, one verdict each

Investor. It earns a second read for one reason: it stakes everything on a single number, states the bar it must clear, and says it will report a null. What it does not have is a market built from the bottom, a buyer whose recurring bill changes, or anything the day after publication that a competitor cannot copy. The two-handset test on shipping recorder apps decides this, and until it comes back there is nothing to price.

Founder. The thesis is intact and the memo undersold it in three places. We conceded all three outputs to pre-installed products that are not established to ship two of them, we conceded the channel question by leaving distribution out entirely, and we conceded the microphone question when one of the three paths that would reopen it shipped in the release we lead on. The screen-lock claim goes the other way and outruns its source, and the breakeven does not reproduce from our own numbers, so both have to be fixed before anyone outside reads this.

Domain generalist. The field will not accept this document as written, because it contains no word error rate, no test set, no speaker-error figure and no latency number of its own, while quoting a rival's latency with its recording length attached. It never says whether the transcript names speakers, which is the first thing a practitioner asks and the one thing neither platform gives a third-party app. Two of its platform facts are stated harder than their sources support, and the headline shipment figure is a tracker number where a fetched one exists.

Speech and voice machine learning. Two of the three mechanisms contradict the project's own registered findings: the session-token wall on the summariser and the suspended app that cannot raise the alarm it exists to raise. The completion bar measures the wrong thing, since a recording can survive to the end and still be unreadable, so it needs a quality floor beside it. The concession about undetectable silence is too generous, and the sessions the memo proposes to collect are the most valuable asset in the plan if they can be collected at all.

Mobile audio platform engineer. The segmented writes are sound engineering and the rest of the capture story is not yet engineering. A killed process warns nobody, so the in-room warning needs a pre-armed notification that a heartbeat cancels; recovery is a third tap and it is unbudgeted. The channel sentence, the screen-lock sentence and the trend-line window are each wrong against a documented interface or a documented release, and the Android test applies the wrong stimulus to the one Android handset it names.

Consumer platform strategist. The memo is disciplined about its unknowns and then drops the store commission out of its only unit-economics line, in a category where that rate is currently unsettled in three jurisdictions. The summarisation half cannot run a meeting in one pass and the document never says which architecture it is building. The convergence answer is contradicted by surfaces both platforms already ship, and the store-review gate that comes before both proposed tests is not mentioned at all. It also gives away the one thing the platform has not taken.

Consumer subscription go-to-market. It retires the hardware buyer, names no replacement, and then prices a generic subscriber, while at most a third of the demand it just measured is reachable in software. There is no channel anywhere, in the category where acquisition cost is what kills companies, and the evidence says the object itself was the channel. It says why we beat the hardware and never why anyone pays us instead of five free products, and the wedge as written is the incumbent's closing argument, not ours.

Recording-consent and privacy counsel. The technical and pricing claims check out; the one legal clause is unsourced and states the consent posture backwards, treating all-party as the exception when a single out-of-state participant makes it the national default. "Put the phone face down" markets covert capture in a category with four pending actions, none of which the memo mentions while naming three of the defendants as free competitors. The dataset the defensibility claim depends on is described in a shape that cannot be collected lawfully, and the lawful shape is already written down. The in-room warning, though, is the cheapest compliance instrument in the product.