

HOW WE TESTED

Best Free AI Note Taker Apps in 2026: Full Scoring & Receipts

Our complete proof pack

This document backs up every claim in our comparison table. We ran the same public clip through all five tools' free plans and graded the results against independently published sources. If you ever doubt a line in our review, this is where you check our work, including links to the original clip and to each tool's actual output.

01 Check it yourself

The clip and the published sources of truth (what was actually said):

- › The recording we tested: [Alphabet Q1 2026 earnings call \(YouTube\)](#)
- › Alphabet's official published remarks: [blog.google, Sundar Pichai Q1 2026 remarks](#)
- › Alphabet Q1 2026 SEC earnings release: [abc.xyz/investor](#)
- › Secondary reference transcript: [Motley Fool transcript](#)

Note: the Motley Fool transcript omits the intro and carries its own typos, so we used it only as a secondary cross-check, never as the standard a tool was judged against.

Each tool's actual output (open and read the real results):

- › Jamie: [Jamie summary & transcript](#)

- › Otter: [Otter output](#)
- › Fireflies: [Fireflies output](#)
- › Fathom: [Fathom output](#)
- › Granola: [Granola notes](#)


02

Methodology and invitation to verify

We have set out our method in full so that any reader, or any company named here, can reproduce this test and check our findings independently. We do not ask anyone to take our conclusions on trust.

All five tools were assessed under identical conditions. We installed the desktop application of each tool on the same Mac computer and captured a single public recording, Alphabet's Q1 2026 earnings call, simultaneously across all five, halting each at approximately the eleven-minute mark (shown below). Every tool therefore received the same audio, at the same moment, in the same environment.

One exception. Granola's scores come from a separate run, noted under the table.



Each tool was then scored out of 100 across three defined categories, against a fixed set of criteria applied uniformly to all:

- **Summary quality:** whether it captured the key result, covered the principal points, reported every figure accurately, and introduced no information that was not stated.
- **Transcription quality:** whether it spelled difficult proper nouns correctly, recorded the figures accurately, and reproduced the words without omission or duplication.
- **Speaker identification:** whether it correctly attributed each statement to the person who made it.

We're Jamie, and Jamie is one of the tools tested here, so we're not asking you to take our word for it. Every tool was graded against independently published sources (Alphabet's official remarks and SEC release), and every claim below is backed by each tool's own output, all linked. Check any line yourself. Each tool was credited for what it demonstrably did well, and every finding is scoped strictly to this recording, on each tool's free plan, on the date stated. (Note: on

free plans some exports are gated, e.g. Fireflies' transcript is viewable but not downloadable without upgrading; see Exhibit C. This did not affect scoring, as the transcript was fully visible.)

We invite you to verify this for yourself. You may install the same tools and trial them on your own meetings, or you may examine our scoring directly: open the recording, copy each tool's transcript, and submit them to Claude, or any independent AI, together with the official transcript, and ask which captured the most and erred the least. We are confident your own assessment will align closely with ours. Our complete reasoning, supported by verbatim evidence for every claim, follows below for line-by-line scrutiny.

03

The scores at a glance

SCORED OUT OF 100	Jamie	Granola	Fireflies	Otter	Fathom
AI summary	90	80	80	84	68
Transcription	90	76	82	72	60
Speaker identification	98	10	28	58	10

Note. Granola was retested on 8 August 2026, on the same clip and the same free plan. On the first run we could not find its transcript, which is why we ran it again. All three Granola scores come from the retest, so screenshots taken during the first session may not match them.

04

Row 1, AI Summary

Jamie, AI Summary

FINDING

Best at: depth plus attribution. In our test, on this clip, Jamie produced the most detailed summary of the five, a short executive-summary block up top,

then a fully expanded version underneath, so you can skim or read the whole record from one file. It's also one of the few here that ties points to a named speaker ("According to Sundar Pichai..."). The trade-off is length: the full summary is deeply nested, so finding a single number can mean scrolling through several indent levels. That's a style call, not a flaw, but worth knowing if you want answers at a glance.

PROOF (VERBATIM)

- Two-tier structure: Jamie output contains "## Executive Summary" and "## Full Summary".
- Speaker attribution: Jamie output, "According to Sundar Pichai, the first quarter of 2026 was highly successful..."
- "Most detailed of the five": Jamie and Granola are the only two summaries carrying all of 16B tokens/min from 10B, both TPU chips split, Vera Rubin NVL72, and Deep Research/MCP; the other three (Otter, Fireflies, Fathom) each omit at least one.
- Deep nesting: Jamie output uses four indent levels (e.g., bullets under "### AI Infrastructure & Model Performance" to TPU8T to detail).

Granola, AI Summary

FINDING

Best at: clean, complete, and easy to read at once. In our test it sorted everything under clear topic headings and captured nearly all the key points, arguably the tidiest single-page summary of the group (our view). Two things to weigh: the notes carry no speaker attribution, and it lightly garbled one product name, writing "VO 3.1 Lite" for Google's Veo 3.1 Lite. The summary also carries no per-line timestamps, and Granola keeps no audio at any point, so there is no recording to jump back to.

PROOF (VERBATIM)

- Topic headings: Granola summary, "Overall Performance", "Google Cloud Highlights", "AI Infrastructure and Models", "Gemini Models and Generative

Media", "Search and Product Updates".

- Both TPU chips split: Granola summary, "TPU 8T: 3x processing power of Ironwood, 2x performance for model training" and "TPU 8i: 80% better performance per dollar for low-latency inference". It also carries "16B tokens per minute (up from 10B last quarter)" and "NVIDIA Vera Rubin NVL72".
- No speaker attribution: Granola output contains no speaker names; bullets are unowned.
- Invented product name: Granola summary, "Imagen 2: reached 1B images in nearly half the time of Imagen 1". The call said Nano Banana 2, and Granola's own transcript reads "Nano Banana two reached 1,000,000,000 images", so the error is in the summary layer. Its June summary named Nano Banana 2 correctly. The summary also writes "VO 3.1 Lite" for Veo 3.1 Lite.
- No timestamps in the summary: Granola output has a date line "Sat, 08 Aug 26" but no per-line times. Its transcript does carry per-line start and end times; what is missing is audio, since Granola stores none.

Fireflies, AI Summary

FINDING

Best at: fast, scannable notes plus a separate action-item list. In our test it auto-extracted to-dos and assigned them to owners, and it timestamped lines so you can jump to a moment in the recording. Good for task-driven meetings and a quick read. Two trade-offs in our test: the notes are shorter and skip some detail the deepest summaries kept (no 16B-tokens figure, no Vera Rubin, no Deep Research/MCP); the summary notes carry no speaker attribution (the bullets aren't tied to who said what), in the transcript Fireflies does label speakers, but maps them unreliably (see Speaker Identification); and the timestamps came out repetitive, almost every line tagged the same time, which weakens their "jump to proof" value.

PROOF (VERBATIM)

- Action items with owners: Fireflies output, "Action items", "Google Cloud Team", "Sundar Pichai / Alphabet Investor Relations".
- Timestamps present and repetitive: Fireflies output tags nearly every notes line "(03:19)".
- Omitted detail: Fireflies output contains no mention of "16 billion tokens," "Vera Rubin," or "Deep Research / MCP," all present in Jamie and Granola outputs.
- No speaker attribution on notes: Fireflies notes bullets are unowned.

Otter, AI Summary

FINDING

Best at: a clean overview plus a labelled outline, with speaker turns and a shareable source link, pleasant to navigate. Importantly, in our test its summary was accurate: the intro it describes (the operator's listen-only housekeeping, the safe-harbor and GAAP/non-GAAP note, the investor-relations website) was genuinely spoken on the call, so it isn't invented. It correctly attributes the welcome to the operator, the host to Jim Friedland, and the CEO to Sundar Pichai. It runs a touch less granular than the deepest summaries, it generalizes "TPU 8" rather than splitting the two chips, but on substance it holds up.

PROOF (VERBATIM)

- Labelled outline with attribution: Otter output, "Speaker 2, Jim Friedland, introduces Sundar Pichai..." and "Speaker 1 welcomes everyone...explains the listen-only mode."
- Intro was actually said (so not invented): recording shows the operator's "all participants are in a listen-only mode" and Jim's "we will present both GAAP and non-GAAP...measures" (present in Jamie, Otter and Fathom transcripts).
- Less granular: Otter output says "introduction of TPU 8" rather than naming TPU 8t / 8i separately.

Fathom, AI Summary

FINDING

Best at: a polished, presentation-ready layout, meeting purpose, clear takeaways, next steps, recording link. It's the nicest of the set to hand to someone (our view), and it got the headline numbers right (Search +19%, Cloud +63% over \$20B, backlog ~\$460B, 40% QoQ, latency >35%, cost >30%). But in our test its summary introduced several specific errors a numbers-critical reader should know about: it stated Gen-AI revenue grew "+799%" when the figure given was "nearly 800%"; it merged the two separate TPU chips into one, attributing both 3x training and 80% inference gains to a single "TPU-8i"; it listed Gemini 3.1 Pro's gains as "reasoning, multimodal understanding, and planning" when the third item said was "cost"; and it named the open model "MF4" instead of Gemma 4. The polish can make confident errors easy to miss, which is why, for exact-figure work, we'd treat it as the highest-risk of the group (our opinion).

PROOF (VERBATIM)

- Headline numbers correct: Fathom output Key Takeaways, "Search revenue up 19% and Cloud revenue up 63%", "backlog nearly doubling QoQ to over \$460B".
- 799% error: Fathom output, "Gen AI Product Revenue: +799% YoY." Actual: Pichai said gen-AI revenue "grew nearly 800% year-over-year" (blog.google / IR, also quoted by TheStreet).
- Chip merge: Fathom output, "TPU-8i: 8th-gen TPU...Training: 3x...Inference: 80%." Actual: two chips, TPU 8t (training, 3x) and TPU 8i (inference, 80%), per Motley Fool news ("TPU 8t and TPU 8i") and the recording.
- Attribute swap: Fathom output, "reasoning, multimodal understanding, and planning." Actual (recording, per Jamie/Otter/Granola transcripts): "reasoning, multimodal understanding, and cost."
- Name error: Fathom output, "MF4: Most intelligent open model." Actual: "Gemma 4, our most intelligent open model" (Jamie/Otter and Fathom's own

transcript all say Gemma 4).

05 Row 2, Transcription

Jamie, Transcription

FINDING

Best at: a clean, ready-to-use transcript that already names the speakers. In our test, Jamie's transcript arrived with each turn labelled by real name, Operator, Jim Friedland, Sundar Pichai, rather than generic "Speaker 1," and timestamped. Every key number was exact. The misses are minor and worth an eyeball: it wrote "Ashkenasi" for the CFO's name (Anat Ashkenazi), "ViO 3.1 Lite" for Veo 3.1 Lite, and "agent decoding" where the audio said "agentic coding." There's also a visible glitch where the legal boilerplate stutters ("non-GAAP, non-GAAP, non-GAAP..."). Numbers stayed perfect; just verify the names.

PROOF (VERBATIM)

- Named speakers + timestamps: Jamie output, "Jim Friedland 00:31 - 01:35", "Sundar Pichai 01:37 - 10:06."
- Numbers exact: Jamie output, "Revenue grew 63%, exceeding \$20 billion...over \$460 billion" matches blog.google verbatim; "16 billion tokens per minute...up from 10 billion," "three times the processing power of Ironwood and two times the performance," "80% better performance per dollar," "nearly 800%," "45%," "40%," "more than 35%," "more than 30%," "70 languages."
- Name slip: Jamie, "Anat Ashkenasi"; correct: Anat Ashkenazi (Alphabet CFO, per SEC release).
- Product slip: Jamie, "ViO 3.1 Lite"; correct: Veo 3.1 Lite (Fireflies output: "Veo 3.1 Lite").
- Audio-contradiction slip: Jamie, "agents, and agent decoding"; audio said "agentic coding" (Otter: "agents and agentic coding"; Fathom: "agents, and

agentic coding").

- Stutter: Jamie output, "non-GAAP, non-GAAP, GAAP and non-GAAP, non-GAAP..."

Granola, Transcription

FINDING

Best at: number fidelity. In our test, Granola's transcript carried every key figure correctly, sixteen out of sixteen and the best of the five. It also spelled Broadcast, Anat Ashkenazi, Ironwood, Vera Rubin and Gemma 4 correctly, matching Fireflies on hard proper nouns. Two things cost it. The proper nouns it did miss, and seven phrases where a word dropped out. Granola keeps a full transcript under every note, opened with the small waveform button beside the chat bar labelled Show transcript, and the whole transcript can be searched and copied from that panel.

PROOF (VERBATIM)

- Numbers correct: Granola transcript, "revenue grew 19%", "Revenue grew 63%, exceeding \$20,000,000,000 for the first time", "over \$460,000,000,000", "16,000,000,000 tokens per minute ... Up from 10,000,000,000", "three fifty million", "70 languages", "grew nearly 800% year over year", "outpaced their initial commitments by 45%".
- One duplicated fragment only: the transcript opens "Thank you for Thank you for standing by", and no whole passage repeats. Counted as repeated eight word runs, Granola has 1, Fireflies and Otter 0, Jamie 4 (the non-GAAP stutter) and Fathom 86.
- Proper nouns correct where others slipped: Granola, "Broadcast" (Fathom: "Broadcast"); "Anat Ashkenazi" (Fathom: "Anant"); "Ironwood"; "NVIDIA Vera Rubin NVL72".
- Name slips: Granola, "Sundar Peccay" (correct: Sundar Pichai); "Philip Schindler" (said Philipp); "TPU AT" (correct: TPU 8T); "VO 3.1 Lite" (correct: Veo 3.1 Lite); "AI Rator portfolio" (correct: AI accelerator portfolio).

- Dropped words: Granola, "we have quickly the Gemini 3.1 series"; "our deep research agent caught a big including MCP support"; "bringing health helpful AI"; "with the that enterprise customers need".
- Transcript is readable and exportable: Granola's help centre, "Copy all ... copy the entire transcript to your clipboard" and "Search ... through your transcript to find specific words or phrases" (docs.granola.ai/help-center/taking-notes/transcription).




Exhibit D. Granola's Quick Note dialog, captured 23 June 2026, states it "will transcribe, but won't know who is speaking" and "will write notes for you once the transcript has stopped." In our June run we did not open the transcript panel.

Fireflies, Transcription

FINDING

Best at: accurate, complete text. In our test, Fireflies' transcript ran clean end to end, no dropped or duplicated passages, and every key number was exact (63%, \$20B, \$460B, 16B from 10B, 3×/2×, 80%, nearly 800%, 45%, 70 languages). It split TPU 8T and TPU 8I correctly and got Veo 3.1 Lite and Anat Ashkenazi right. Minor mishears: "Philip Schindler" (said Philipp), "brand cast" (Broadcast), "AIO views" (AI Overviews), "Liria 3" (Lyria). Turns are tagged Speaker 1 to 4 with no names in the body, see Speaker ID. On the free plan the transcript is viewable in-app but not downloadable, Download → Transcript routes to an upgrade wall (see Exhibit C), so exporting it requires a paid plan.

PROOF (VERBATIM)

- › Numbers exact: "Revenue grew 63%, exceeding \$20 billion... over \$460 billion"; "16 billion tokens per minute... up from 10 billion"; "nearly 800%."
- › Hard names: "TPU 8T," "TPU 8I," "Veo 3.1 Lite," "Anat Ashkenazi."
- › Complete, no duplication.
- › Slips: "Philip" (Philipp); "brand cast" (Brandcast, per blog.google).
- › Free-plan export wall: Download → Transcript shows PDF/DOCX/SRT/CSV/JSON/MD, but the Download action prompts "Upgrade to download meeting" (Exhibit C). Transcript remains viewable in-app.




Exhibit C. On the free plan, Fireflies' transcript is viewable in-app but cannot be exported. Download → Transcript offers PDF/DOCX/SRT/CSV/JSON/MD, but the Download button routes to an "Upgrade to download meeting" wall. We scored the viewable transcript; export requires a paid plan.

Otter, Transcription

FINDING

Best at: a complete, faithful transcript end to end. In our test it even captured the opening housekeeping and safe-harbor section that the secondary Fool transcript omits, and that content was genuinely spoken, so it's a point in Otter's favour. The numbers are substantively all there, and it spelled the CFO's name correctly ("Anat Ashkenazi"). The catch is speaker work: the transcript body carries no names, Otter reports voices by share of talk-time ("Speaker 1 (84%)...") and leaves the text unlabelled, so identifying who's who is manual. Two accuracy slips in our test: it garbled the chip names to "TPU 80" and "TPU AI" (for TPU 8t and TPU 8i), and wrote "three 50 million" for 350 million.

- Captured the (real) intro: Otter output, "all participants are in the listen-only mode...press star one"; corroborated as spoken by Jamie/Fathom transcripts.
- Correct CFO name: Otter output, "Anat Ashkenazi."
- No names in transcript body; talk-time labels only: Otter output, "Speakers Speaker 1 (84%), Speaker 2 (11%), Speaker 3 (4%), Speaker 4 (<1%)" with unlabelled body text.
- Chip garbles: Otter, "TPU 80 provides high performance model training" and "TPU AI delivers cost-effective low latency inference"; correct: TPU 8t / TPU 8i (Motley Fool news: "TPU 8t and TPU 8i").
- Number glitch: Otter, "reached three 50 million"; correct: 350 million (blog.google context; Jamie: "reached 350 million").

Fathom, Transcription

FINDING

Best at: number fidelity. In our test every figure in Fathom's transcript was correct, including the "nearly 800%" it later misstated as "799%" in its own summary. The weaknesses are structural. It placed the operator, the IR host and the CEO under a single "Speaker 2" block, so the transcript reads as if one person delivered the whole opening. This is expected: per Fathom's own help center, it "cannot distinguish between multiple speakers in the same physical room" and "all audio will be attributed to the Fathom user's name in the transcript and summary" (help.fathom.video/en/articles/5500225), so with our single-source playback, who-said-what is effectively lost here (see Exhibit B). It also duplicated several passages verbatim (the 63%/\$460B revenue lines, the Gemini Enterprise block, the generative-media block), and in one duplicated pass dropped a word ("YouTube and Google" instead of "YouTube and Google One"). A few names slipped: "Anant" for Anat, "Broadcast" for Brandcast, and "Liria" for Lyria. It's the messiest record of the set in our test (our characterization).

- Numbers correct (incl. 800%): Fathom transcript, "grew nearly 800% year over year" (matches the call; contrast its own summary "+799%").
- Single Speaker 2 block: Fathom output, "0:00 - Speaker 2" spanning the operator, Jim and Sundar with no further speaker change.
- Duplication: Fathom output repeats "Revenue grew 63%, exceeding \$20 billion...over \$460 billion" and the "Gemini Enterprise is seeing tremendous momentum..." block.
- Dropped word in duplicate: first pass "YouTube and Google One being the key drivers"; duplicate pass "YouTube and Google being the key drivers."
- Name slips: Fathom, "Anant Ashkenazi" (correct: Anat Ashkenazi, SEC release); "I/O, Broadcast, and GML" (correct: Brandcast, blog.google: "I/O, Brandcast and Google Marketing Live"); "Liria 3" (correct: Lyria 3).

⚠ Please note:

- Fathom cannot distinguish between multiple speakers in the same physical room. All audio will be attributed to the **Fathom user's name** in the transcript and summary.
- This workaround is **not ideal for large in-person meetings or conferences**.

Exhibit B. Fathom's own help center confirms it "cannot distinguish between multiple speakers in the same physical room," attributing all audio to the Fathom user, help.fathom.video/en/articles/5500225.

06 Row 3, Speaker Identification

Jamie, Speaker Identification

FINDING

Best at: speaker labels you don't have to fix. In our test, Jamie was the only tool of the five to separate every voice and attach the correct real name to each,

Operator, then Jim Friedland, then Sundar Pichai, with a time range per turn and the hand-offs in the right order. "Operator" stays generic, which is correct, since that speaker is anonymous. The only blemish is a spelling slip inside Jim's speech ("Ashkenasi" for Ashkenazi), a spelling issue, not a mis-attribution; the speaker tags themselves are right.

PROOF (VERBATIM)

- Named, separated, time-ranged: Jamie output, "Operator 00:05 - 00:29", "Jim Friedland 00:31 - 01:35", "Sundar Pichai 01:37 - 10:06."
- Only tool to do so: Otter labels by number only ("Speaker 1 (84%)..."); Fathom uses one "Speaker 2"; Fireflies splits into Speaker 1 to 4 but maps them unreliably; Granola places the operator, Jim Friedland and Sundar Pichai under one "Them" label with no names.
- The spelling slip is inside speech, not a tag: Jamie body, "Sundar Pichai, Philipp Schindler, and Anat Ashkenasi."

Granola, Speaker Identification

FINDING

In our test, Granola placed the operator, Jim Friedland and Sundar Pichai under one "Them" label, with no names and no turn breaks, so who said what is lost in this capture. That matches Granola's own help centre, which states that where speaker tags are not available you see "Me" and "Them", corresponding to your microphone input and your system audio. Credit where it is due: Granola pulls real participant names from the meeting platform on Google Meet through its browser extension and on Zoom on macOS through Settings, and the iPhone app separates speakers as Speaker A and Speaker B. Our clip played through system audio on the desktop, so none of those routes applied here.

PROOF (VERBATIM)

- One label for three people: Granola transcript, a single "Them:" block carrying the operator's welcome, Jim Friedland's safe harbor reading and Sundar Pichai's remarks, with no further speaker change.

- No names anywhere: the Granola transcript header lists only "Meeting participants: Sanduni Yureka"; no speaker in the body is named.
- This is documented behaviour, not a fault of the capture: Granola's help centre, "you'll see Me and Them in the transcript, corresponding to your microphone input and your system audio" (docs.granola.ai/help-center/taking-notes/transcription).
- Named tags exist on other routes: Granola's help centre, "Google Meet ... available on macOS and Windows through the Granola browser extension", "Zoom ... available on macOS through Settings > Preferences", and "Granola uses the display name your meeting platform shows for each participant" (docs.granola.ai/help-center/taking-notes/speaker-attribution).

Fireflies, Speaker Identification

FINDING

In our test, Fireflies split the call into four numbered speakers but mapped them unreliably, it chopped the operator's intro between "Speaker 1" and "Speaker 2," gave Jim Friedland and Sundar Pichai the same "Speaker 3" label (two people read as one), and split one of Jim's sentences off as "Speaker 4." No real names appear; attribution can't be trusted as-is.

PROOF (VERBATIM)

- Operator split: "Speaker 1: 00:00 Welcome, everyone" then "Speaker 2: 01:44..."
- Two merged: Jim "Speaker 3: 02:09..." and Sundar "Speaker 3: 03:15..."
- No real names: only "Speaker 1 to 4."

Otter, Speaker Identification

FINDING

Best at: telling you who the voices were, in the summary. In our test Otter's outline correctly identified the operator's welcome, Jim Friedland as host, and Sundar Pichai as CEO, so you're not left guessing. The limitation is in the

transcript itself: it has no name tags, just paragraph breaks where the speaker changes, so building a labelled record means identifying each voice yourself. One detail worth noting: Otter reported four distinct voices on a clip where only three people actually spoke, so its voice-splitting was slightly over-counted here.

PROOF (VERBATIM)

- Names worked out in the outline: Otter output, "Speaker 2, Jim Friedland, introduces Sundar Pichai..." and "Speaker 1 welcomes everyone..."
- Transcript body unlabelled: Otter body text runs in unnamed paragraphs.
- Four counted, three spoke: Otter output, "Speaker 1...Speaker 2...Speaker 3...Speaker 4 (<1%)"; the clip has three speakers (Operator, Jim, Sundar), as the named transcripts show.

Fathom, Speaker Identification

FINDING

In our test, Fathom applied a single speaker label and nothing more: the operator, the IR host and the CEO were all placed under one "Speaker 2," so the transcript reads as if one person delivered the entire opening, including the operator handing to Jim, and Jim handing to Sundar. No real names appear, and the distinct hand-offs are lost, which is why speaker identification scored at the bottom here. It did at least apply a label, but for "who said what," this output doesn't deliver it (our assessment).

PROOF (VERBATIM)

- Single label over multiple speakers: Fathom output, "0:00 - Speaker 2" covering the operator's "press star one", Jim's safe-harbor, and Sundar's "Thanks, Jim" with no speaker change.
- No real names: Fathom output uses only "Speaker 2."

J Jamie *Tested by the team at Jamie on 23 June 2026, with Granola retested afterwards as noted under the scores. Every quote above is drawn from each tool's own output or the official published sources linked at the top. Scores are scoped to this clip, on each tool's free plan, on the date tested.*

