



*A whisper, when it matters.*

*Notetaking is solved.  
The call isn't.*

A desktop tool that helps sales reps sell more, by telling them what to say in the moment.

# *Every interview is a sales call.*

*And every tool lives on the wrong side of the moment.*

## THE LISTENER

### *Notetakers*

Otter, Fireflies, Gong. They record, transcribe, and tell you exactly what happened. After.

## THE SPEAKER

### *Nothing, live*

No one helps you in the moment you're actually in, when it could still change the outcome.

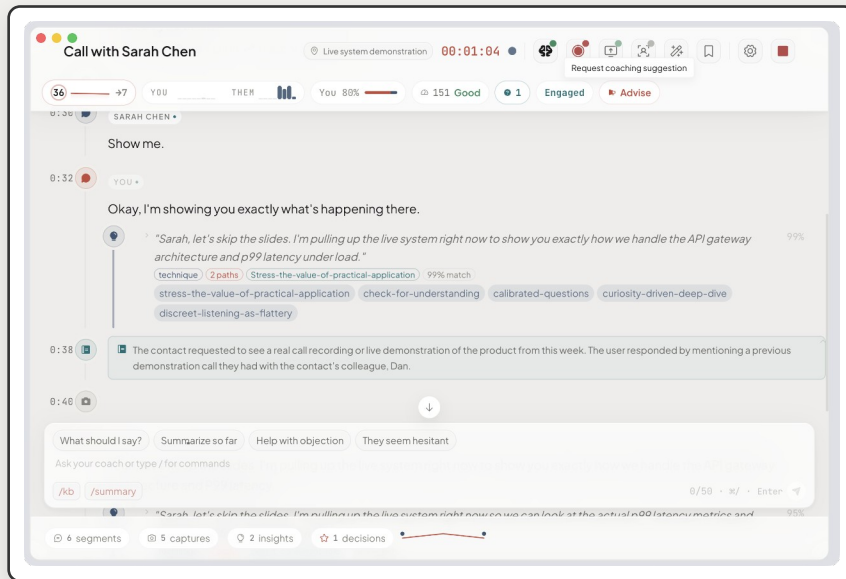
# *A whisper, when it matters.*

**For the speaker, not the listener.**

It tells you what to say, in the moment.

**On-device.**

Audio never leaves your Mac. No bot joins the call.



Live coaching overlay, during the call

# The call you'd normally fumble.

Prep in 30 seconds. A whisper at the turn. A debrief after.

**Sarah Chen**  
PRE-CALL PREPARATION

Last spoke 28/05/2026. Main topic: Quick 18-minute alignment with Sarah ahead of the POC kickoff. Watch for API gateway architecture. Watch for API gateway architecture. p99 latency optimization, team scaling past 50 engineers, build vs buy decisions, developer experience metrics.

**WATCH FOR:** 1 default to talking when nervous about a stalled deal. Average talk ratio on first-touch calls: 41%. On stalled-deal recovery calls: 57%. Higher talk = lower rapport. The exit move from a stalled deal is a question, not a presentation. - TRY: When the CISO joins (James + Dr. Patel pattern), pre-send a 1-page architecture-and-audit-trail brief 24h before — let them read it, not present it. Then the call is Q&A on their margin, not pitch on mine. Test on next 5 CISO bring-ins.

**DOCUMENTS**  
Save to your knowledge base  
Session only — available for this call

Upload deck/proposal

**GOAL**  
Close the deal

1 • Prep (30s)

Call with Sarah Chen  
00:01:04  
request coaching suggestion

56 47 101 118  
Your 80% 151 Good 1 Engaged Admin

**SHOW ME.**

0:32 Sarah Chen  
Okay, I'm showing you exactly what's happening there.

0:34 The contact requested to see a real call recording or live demonstration of the product from this week. The user responded by mentioning a previous demonstration call they had with the contact's colleague, Dan.

0:44

What should I say? Surprise so far Help with objection They seem hesitant  
Ask your coach or agent for commands

1 segments 5 captions 2 insights 1 decision

2 • Live whisper

**Elena Rodriguez**  
Head of Product, CloudScale Systems  
May 27, 2026 - 50:08 min - 14th session  
Neutral  
68  
RAPPORT  
Good Rapport Established

**INTELLIGENCE BRIEF**

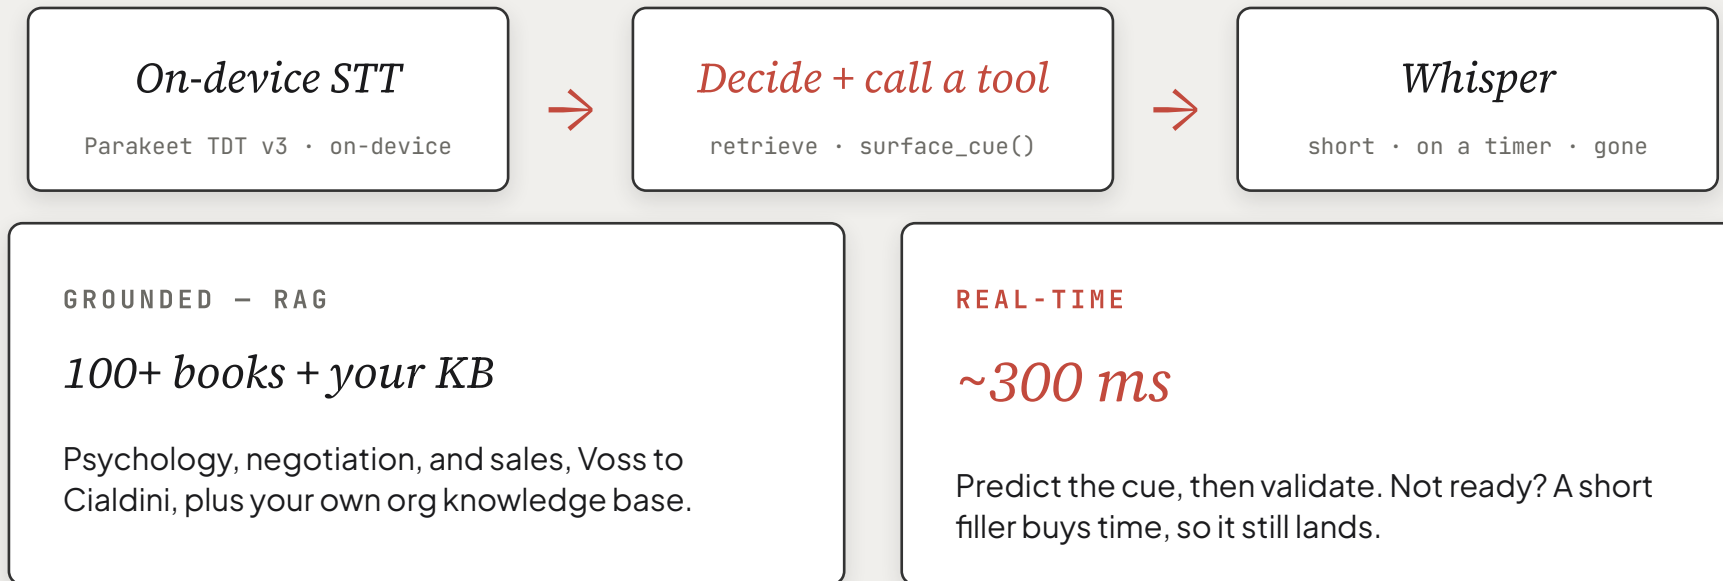
**Confirm the launch date alignment with their Q3 release**  
50-minute co-design working session with Elena at CloudScale Systems. Goal: Confirm the launch date alignment with their Q3 release. Closing rapport 68/100. Technical concerns surfaced but addressable. Need to win Kai independently. The next touch needs to address the surfaced concern explicitly before we attempt to advance.

68 Rapport  
0% Talk Balance -Them-heavy  
50m Duration

KEY MOMENTS ACTIONS & FOLLOW-UPS

3 • Debrief

# *An LLM that decides when to call a tool.*



# *Four layers. All sliding to real-time.*

**Notetakers** — Otter, Granola, Fireflies

commoditized

**Revenue intel** — Gong, Clari

for your manager

**Voice agents** — Vapi, Retell

AI is the caller

**Real-time assist** — coach a human, live

**OPEN**

# *Three unsolved problems in a trenchcoat.*

## LATENCY

### *300ms vs 1.5s*

The target is 300 milliseconds. Production voice runs at a second and a half. The demo-to-real gap is the whole game.

## TURN-TAKING

### *When to shut up*

The hard part isn't generating a cue. It's not firing mid-sentence. Naive barge-in is about half wrong.

## EVAL

### *Was it helpful?*

No benchmark for cue quality. I built eleven eval configs, judge models plus humans. Honestly, the weakest part.

# *Real intent, not vanity metrics.*

**8**

signed LOIs

**~30**

waitlist

**483**

test files

**11**

eval configs

**Solo · 3 months · zero paid marketing**

Parakeet TDT v3 · on-device + bring-your-own LLM · RAG · native macOS

*Only text ever leaves the device, never your audio.*

# *Notetaking is solved. The call isn't.*

*The right ten words, whispered in time.*

## TRY IT YOURSELF

Reach out and I'll send you the  
download link.

macOS · Apple Silicon (M-series)

## BUILDING IN PUBLIC

Got ideas or feedback? Let's  
brainstorm.

Tell me what I'm getting wrong.

# Where Sussur sits.



# *Built with.*

## SPEECH-TO-TEXT

Parakeet TDT v3, running on-device

## REASONING

Bring-your-own LLM, OpenRouter / OpenAI / Ollama

## KNOWLEDGE

RAG over 100+ books (psychology, negotiation, sales) + your org KB

## REAL-TIME

Predict, validate, filler pipeline, ~300 ms budget

## EVALS

LLM-as-judge + human ratings, 11 configs

## APP

Native macOS, Apple Silicon

# *Hitting ~300 milliseconds.*

## THE BUDGET

*~300 ms, end to end*

STT, decide, retrieve, render. A cue two seconds late is just noise, and production voice systems run at a second and a half.

## THE TRICK

*Predict, then validate*

Speculate the cue and check it. If it's not ready, speak a short filler while it finishes, so you never feel the lag.

# Three forces, 2026.

## ON-DEVICE

### *Good enough*

Parakeet TDT v3 and local LLMs now run real-time on a Mac. No cloud, no bot in the room.

## LATENCY

### *Crossed the line*

Live coaching went from 2–4 seconds to ~300 ms. Usable in the moment for the first time.

## THE GAP

### *Wide open*

Notetaking is commoditized and platform-owned. The individual seller still has no live coach.

## *Two SKUs. Over 90% margin.*

### PERSONAL

*€49–149 one-time*

Bring your own LLM key (OpenRouter, OpenAI, Ollama). Zero variable cost to us. Sold on a credit card, before procurement.

### BUSINESS

*€30–120 / month*

Managed inference. About \$0.0025 per coached minute, 91–96% gross margin. Compounds as reps bring teams.

# *On-device is the moat.*

## **Audio never leaves the Mac.**

Transcription is local; only text goes up.

## **Recording is now a legal risk.**

Otter faces a class action that alleges it secretly recorded calls (Aug 2025).

## **11 US states require all-party consent.**

On-device sidesteps the bot-in-the-room problem.

