

Native enterprise AI provides formulaic filler, not actual insight.

**Zoom & Teams
Native AI:**
Produces generic,
repetitive summaries.



**Google Workspace
(Gemini):**
Bundled note-taker
is welded to Meet;
dead weight if you
don't use it.

The Core Problem:
You need a real note-taker, not a
bot that spits out empty templates.

The entire transcription stack is just three commoditized open-source blocks.

Audio Capture:
macOS pathways like
ScreenCaptureKit or
CATapDescription.
Captures system and mic.

Speech-to-Text (STT):
The heavy lifters running
on Apple Neural Engine
(Whisper, Parakeet, Qwen3).

Summarization:
Local LLMs (Ollama) or
BYOK for Claude/GPT.

Optional 4th block:
Speaker Diarization
(FluidAudio) to cluster
who said what.

Hand-rolled macOS audio routing is a steep wall for users

Step 1 (The Block):

macOS privacy aggressively blocks direct system audio capture.

Wa-!

Step 2 (The Hack):
Install BlackHole
(virtual audio driver).

Step 3 (The Headache):

Open Audio MIDI Setup -> Create a Multi-Output Device -> Route system audio through it.

Case Study (Notes4Me):

An 8-star GitHub project that proves the concept but exposes every layer of this frustrating stack. It's a working minimum loop, not a usable product.

The polished SaaS tier sells UX, but the price is friction and privacy.





Choose Your Fighter

Pick 1:
Meetily
(11.6k stars)

Profile: The Install & Go
Granola alternative.

Stack: Rust + Tauri.
GPU acceleration.

Why it wins:

- ★ Zero Xcode or BlackHole required. Just works out of the box.
- ★ Free community edition.

Pick 2:
pasrom/meeting-transcriber
(19 stars)

Profile: The Mac Purist &
Bilingual Choice.

Stack: Native Swift,
CATapDescription (no drivers).

Why it wins:

- ★ Runs WhisperKit or Qwen3-ASR locally on Apple Silicon.
- ★ Flawless mix of English and Chinese transcription.

The closed-source window is shutting. Subscription fees are just friction fees.