

PROJECT BRIEF



PREPARED

July 2026

AUTHOR

Jordan Bee

Wispr Lite — Fully-Local Voice Dictation

A from-scratch, fully-local clone of Wispr Flow: hold a key, speak, and cleaned-up text pastes at your cursor in any app — with the audio never leaving the machine.

SIGNATURE CAPABILITY

Replaced a paid dictation subscription with a local build where audio never leaves the Mac — speech is transcribed on-device and only optional text cleanup touches the network.

OVERVIEW

Wispr Lite is a menu-bar dictation tool I built to replace my paid Wispr Flow subscription. Holding Right Option is push-to-talk, double-tap locks recording on, and a tap stops it. The mic streams in ~6-second chunks to a local Whisper model (mlx-whisper large-v3-turbo) running on-device, so audio never leaves the machine; only the transcript optionally goes to a light LLM for cleanup tuned to whatever app is frontmost. It reads Wispr Flow's own dictionary live from its SQLite store plus my own vocab file so names and jargon come out right, then pastes the result at the cursor. Very short dictations skip the cleanup step for an instant paste.

TECHNOLOGY

Python 3.14, mlx-whisper (large-v3-turbo, on-device Apple-silicon STT), Gemini 2.5 Flash-Lite (optional text cleanup), rumps (menu bar), sounddevice / PortAudio, pynput hotkeys, SQLite (reads Wispr Flow's dictionary), macOS Accessibility + Microphone APIs.

HIGHLIGHTS

- ◆ **On-device speech-to-text** (mlx-whisper large-v3-turbo) — audio never leaves the Mac; only optional text cleanup touches the network.
- ◆ **Streaming transcription** in ~6s chunks with a forced flush at 15s, so long dictations stay responsive.

- ◆ One-key control on **Right Option**: hold = push-to-talk, double-tap = lock on, tap = stop.
- ◆ **Dead-mic detection**: probes RMS per input device and skips silent inputs (e.g. an unplugged USB interface) that make Whisper hallucinate.
- ◆ Self-heals PortAudio device changes by **os.execv restarting itself** instead of crashing on a segfault-prone reinit.
- ◆ Reuses **Wispr Flow's own dictionary** (live from its SQLite) plus a personal vocab file; a fast path skips LLM cleanup for 3-word dictations.

STATUS

Live and in daily use, replacing the paid Wispr Flow subscription; a single ~600-line Python app on the Mac.



Jordan Bee
jordan.bee2012@gmail.com
+44 7532 722082