

LiveKit Meeting Note-Taker — Repo Cheat Sheet

Fora Soft Learn · AI for Video Engineering

1 - The repo in six files (Apache-2.0, livekit-agents 1.x)

src/agent.py	Entrypoint: AgentServer + dispatch + wiring + shutdown summary
src/transcriber.py	One STT session per participant (multi-user transcriber)
src/notes.py	TranscriptStore + one-pass LLM summary + .md / .json writers
src/recording.py	Optional LiveKit Egress room recording (off by default)
src/token_server.py	FastAPI: explicit dispatch + mint a user join token
tests/test_notes.py	Unit tests for the notes layer (no network needed)

2 - Quickstart (about five minutes)

```

uv sync                                # install dependencies
cp .env.example .env.local            # add LIVEKIT / DEEPGRAM / OPENAI keys
uv run python src/agent.py download-files # fetch VAD / turn-detector weights once
uv run python src/agent.py console      # talk to it in your terminal
uv run python src/agent.py dev          # run as a worker vs your server

```

3 - Why it is listen-only (and cheaper)

```

raise StopResponse() + audio_output=False  -> the agent never speaks.
One STT session per participant  -> attributed lines, no diarization step.
No text-to-speech  -> the heaviest cost layer is gone.

```

4 - Cost shape (LiveKit Cloud, listen-only)


Billed: WebRTC participant minutes + agent session minutes + downstream GB
 Plus: speech-to-text per minute (~\$0.006/min). No TTS line item.
 Re-verify tiers at livekit.io/pricing; full worked example in the lesson.

5 - Build-vs-buy checklist

- ☐ Notetaker lives inside your own app? Build on this repo (Cloud or self-host).
- ☐ Strict privacy or data residency? Self-host the open-source LiveKit server.
- ☐ Just need a transcript of a Zoom call? Buy a meeting bot (e.g. Recall.ai ~\$0.50/hr).
- ☐ Need the agent to speak? Add a TTS stage - and expect the cost to roughly double.
- ☐ Recording on Cloud failing? Use `file_outputs=[...]`, not a single 'output' field.

Companion to: 'LiveKit Meeting Note-Taker — Full Repo'

www.forasoft.com/learn/ai-for-video-engineering/articles-ai/livekit-meeting-notetaker-full-repo



GO DEEPER · VOLUME 7 OF 7

AI for Video: A Complete Guide to Intelligence in Video Products

From Real-Time Pipelines and Agents to Serving, Cost, Compliance, and Verticals

By Nikolay Sapunov, CEO at Fora Soft · 1,221 pages · Kindle & paperback on Amazon

[Read it on Amazon →](#)
amazon.com/dp/B0GX3DGBH4

This cheat sheet maps the six files of the note-taker repo, the quickstart and the listen-only cost shape. Volume 7 explains what sits behind it: AI inside the WebRTC pipeline, live captions, agent primitives and memory, serving infrastructure, cost per stream-hour, and compliance for recording and transcription.

WHAT YOU WILL FIND INSIDE

- AI inside the WebRTC pipeline
- Agent primitives, memory and tools
- Live captions and transcripts
- Compliance for recording and transcripts

THE FULL HANDBOOK · 7 VOLUMES · 9,000+ PAGES · ALL ON AMAZON

Vol. 1 · Video Encoding — pixels and color to H.264, HEVC, AV1, quality metrics and AI

Vol. 2 · Video Streaming I — network transport, contribution and ingest, HLS, MPEG-DASH, CMAF

Vol. 3 · Video Streaming II — ABR, CDNs, players, WebRTC, DRM and the economics of delivery

Vol. 4 · Audio for Video I — PCM, codecs, loudness, accessibility, real-time voice processing

Vol. 5 · Audio for Video II — jitter buffers, lip sync, Atmos, quality metrics, neural codecs

Vol. 6 · AI for Video I — models and inference, computer vision, speech, search, generative video

Vol. 7 · AI for Video II — real-time pipelines, agents, serving, cost, compliance, verticals

Want a meeting note-taker inside your own product instead of a third-party bot?

We will adapt the note-taker to your app, models and storage and get it running on your cloud or LiveKit. 250+ video products shipped since 2005.

[Book a 30-min call →](#)
[WhatsApp →](#)
[Email us →](#)