

# Noise Suppression Selection Checklist

2026 latency + cost + licence comparison, the G.114 budget, and a pre-flight before you ship

## 1 · The 2026 noise-suppression options at a glance

| Suppressor              | Latency | Compute target              | Quality             | Licence      | Best for                 |
|-------------------------|---------|-----------------------------|---------------------|--------------|--------------------------|
| <b>RNNoise</b>          | ~10 ms  | 1 CPU core / browser WASM   | Good (steady noise) | BSD          | Tight compute, browser   |
| <b>DeepFilterNet 3</b>  | ~40 ms  | Modern CPU; RTF 0.19 laptop | Best open quality   | MIT / Apache | Highest open quality     |
| <b>Krisp SDK</b>        | ~25 ms  | All major platforms         | Production-grade    | Commercial   | Cross-platform, zero ops |
| <b>Browser built-in</b> | low     | Browser (free baseline)     | Basic               | Free         | Try this first           |


Latency = added algorithmic latency from each project's docs/papers (May 2026). All three clear a live-call budget; the deciding factors are quality, compute, and licensing. Always test on YOUR noise.

## 2 · The latency budget — fit it under ITU-T G.114 ( $\leq 150$ ms one-way)

| Pipeline stage                                            | Adds             | Running total vs 150 ms  |
|-----------------------------------------------------------|------------------|--------------------------|
| <b>Network + capture/encode/decode (typical baseline)</b> | $\approx 110$ ms | —                        |
| <b>+ RNNoise</b>                                          | + 10 ms          | = 120 ms ✓ room to spare |
| <b>+ Krisp</b>                                            | + 25 ms          | = 135 ms ✓               |
| <b>+ DeepFilterNet 3</b>                                  | + 40 ms          | = 150 ms — at the line   |

## 3 · Pre-flight before you ship a noise-suppression feature

- ☐ Tried the browser built-in first: noiseSuppression: true on getUserMedia — if complaints stop, ship it and stop.
- ☐ Ran the G.114 latency arithmetic for your pipeline; confirmed your remaining headroom fits the suppressor you chose.
- ☐ Confirmed added latency, not demo quality: RNNoise ~10 ms, Krisp ~25 ms, DeepFilterNet ~40 ms.
- ☐ Matched the engine to hardware: one CPU core / phone / browser → RNNoise; capable CPU → DeepFilterNet; everywhere → Krisp.
- ☐ Decided self-host vs buy: free permissive licence (you operate) vs commercial SDK (vendor operates).
- ☐ Checked the privacy constraint: all three run on-device — verified your deployment keeps audio local.
- ☐ Picked the licence deliberately: RNNoise BSD, DeepFilterNet MIT/Apache, Krisp commercial.
- ☐ Measured quality with P.835 / DNSMOS, scoring SIG, BAK, and OVRL separately — not a single before/after clip.
- ☐ Tested on your real noise (home office, ward, classroom), not the vendor's demo recording.
- ☐ Tuned for overall quality, not maximum suppression — aggressive settings damage the voice (low SIG).
- ☐ Avoided double-processing: did not stack the browser suppressor and a custom model on the same stream.
- ☐ Confirmed the real-time factor  $< 1$  on your target device, with headroom for other work on the same core.



GO DEEPER · VOLUME 4 OF 7

## Audio for Video: A Complete Guide to Sound in Video Products

From PCM and Codecs to Loudness, Accessibility, and Real-Time Voice Processing

By Nikolay Sapunov, CEO at Fora Soft · 1,229 pages · Kindle & paperback on Amazon

[Read it on Amazon →](#)
[amazon.com/dp/B0HH6PRZH1](https://amazon.com/dp/B0HH6PRZH1)

This checklist compares RNNoise, DeepFilterNet and Krisp on latency, cost and licence and runs the G.114 budget before you ship. Volume 4 covers the voice chain in depth: noise suppression, echo cancellation, AGC and VAD in real time, speech codecs, and loudness.

### WHAT YOU WILL FIND INSIDE

- Where suppression sits in the voice-processing chain
- Voice activity detection and its latency
- Echo cancellation and AGC alongside suppression
- Speech codecs and loudness after cleanup

### 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

## Need to ship noise suppression without blowing the latency budget?

We will measure added latency and real-time factor on your target devices and pick the engine and licence that fit your product. 250+ video products shipped since 2005.

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