


Quick guide: check a set of recordings in your browser

AudioLab Voice AI QC · one page · 20 October 2026


The page measures every file for loudness, true peak, noise floor, bandwidth, clipping and gaps, and gives one verdict per file with the rule and the number behind it. **You** decide the limits; the page only judges against them.

Five steps

1. Open audiolab.tools/voice-ai-check/qc. Measuring happens in your browser.
2. Drop your files on the page, or press **Load the three as a set** to use the built-in sample recordings. Open the **Dataset** tab: one row per file, with its verdict.
3. Set your own limits: press **Open in grid** on any file, then **Rules** and **Edit the limits**. The page starts with a placeholder (the range of 48 recordings of 3 synthetic voices) and says so in a notice: real recordings usually fail it on loudness or bandwidth, which says they differ from those voices, not that they are broken. Replace it with the range of about twenty of your own good files, the 5th to 95th percentile (by construction about 1 in 10 of your good files will still fall outside it), or pick a published target such as EBU R128.
4. Read the verdicts: **Pass**, **Check** (too close to call) or **Fail**. **Rules that fail most** shows where a batch has a problem; **Open in grid** shows six lanes on one timeline, each with its status (validated or proxy).
5. Press **Save these limits**, or download them as a gate file for your pipeline, and reuse them for every batch.



Dataset tab with the starting limits and the notice above the table.



Rules tab: type a minimum and a maximum per measure; the verdicts follow.

Where your audio goes

This browser check: your audio is measured in your browser and is not uploaded; we never receive the audio. If you sign in and save a run, we keep its numbers, your limits and the verdict, never the audio. **The API, including the audit script for whole folders:** the audio is processed in the request and an uploaded file is deleted within an hour at the latest; we do not use it to train models or improve the engine, and we do not sell or share it. All routes, route by route: audiolab.tools/data-routes

What it does not tell you

- Not the language, the words, who is speaking or whether the transcript is right.
- Whether a file was ever lossy, upsampled or spliced is not judged in the browser: it resamples the file before measuring, which hides those traces. The dataset audit reads each file at its own rate through the API (20 files per call, files uploaded for the analysis; speech and phone codecs stay invisible there too): audiolab.tools/docs/start/audit-a-conversation-dataset
- Bandwidth has no published margin, and the browser reads it differently from the file's own rate (up to about 2 kHz in our test): a reading near your limit is the least certain call. Every measure shows its status (validated, proxy) on the page.

Questions or a measure that looks wrong to you: reply to Nathan on LinkedIn or write to hello@audiolab.tools. If a number is wrong, it goes on our page of where we were wrong, with your name if you like.