# Low-Latency Audio — Part-Sizing Cheat Sheet

*Fora Soft Learn*

The frame-alignment math and the per-codec latency floor for LL-HLS, LL-DASH and chunked CMAF.

## Audio frame math (at 48 kHz)

| CODEC | FRAME | DURATION | INTEGER-SEGMENT HINT |
|---|---|---|---|
| AAC-LC | 1024 samples | **21.33 ms** | 8 s = 375 frames (exact) |
| HE-AAC | 2048 samples | **42.67 ms** | 8 s = 187.5 — use 4 s grid |
| Opus | 20 ms default | **20.00 ms** | any 20 ms multiple is exact |
| Opus (low delay) | 2.5–10 ms | **2.5–10 ms** | smaller floor, more overhead |

## The audio floor you cannot tune away

| TERM | TYPICAL | NOTE |
|---|---|---|
| AAC encoder priming | **~44 ms** | 2112 samples ÷ 48000; discarded on decode, still in the path |
| Frame-grid rounding | **up to ~21 ms** | one whole AAC frame of slack at the part boundary |
| Player playout buffer | **~150–400 ms** | main stability-vs-latency knob; the only big one you tune |

## Authoring checklist — tick before you ship

- [ ] Pick a segment duration holding a whole number of audio frames (8 s = 375 AAC frames at 48 kHz).
- [ ] Set EXT-X-PART-INF PART-TARGET so each LL-HLS part holds an integer frame count, or accept slight DURATION variance.
- [ ] In LL-DASH, let the audio AdaptationSet use a more conservative availabilityTimeOffset than video.
- [ ] Align audio and video presentation timestamps; do not chase the audio floor by shrinking parts.
- [ ] Encrypt / package CMAF chunks once so they serve both LL-HLS parts and LL-DASH segments.
- [ ] Prefer Opus over AAC when you control the player and need a lower latency floor.

### The one rule

You cannot subdivide a 21.33 ms AAC frame. Align timestamps and let audio run on a coarser grid; smaller video parts never shrink the audio floor — they only make it more visible.