

When audio needs an adaptive bitrate ladder, when one rendition is enough, and the real platform numbers.

The arithmetic that decides it

audio share of stream = audio bitrate / (video bitrate + audio bitrate)

Stereo on 1080p: $128 / (6000 + 128) = 2.1\%$ -> **adapting saves nothing -> ONE rendition**

5.1 on mid rung: $640 / (1500 + 640) = 30\%$ -> **a real chunk -> build a ladder**

Decision rules

- No video (music, podcast, radio)? Build a 2-4 rung audio-only ladder (Case 2).
- Surround or immersive audio? Short ladder + a stereo floor for weak links (Case 1).
- Stereo on video, lowest video rung below ~600 kbit/s? Add one lower stereo rung.
- Stereo on video, lowest video rung above ~1 Mbit/s? Ship ONE stereo rendition.
- Never copy the video rung count onto audio — quality plateaus, bits are wasted.

What the platforms actually ship (2026)

PLATFORM	CONTENT	AUDIO RUNGS	CODEC(S)
Netflix	5.1 surround	192 -> 640 kbit/s (+stereo floor)	DD+, AC-3, AAC
Netflix	Dolby Atmos	up to 768 kbit/s	DD+ JOC
Apple HLS	stereo video	one rendition, 32-160 kbit/s	AAC-LC
Apple HLS	5.1 surround	DD 384 or DD+ 192 kbit/s	AC-3, E-AC-3
YouTube Music	audio-only	48 / 128 / 256 kbit/s	AAC, Opus

The one rule

Audio needs a ladder in exactly two cases: surround/immersive sound, and audio-only services. Everything else is one rendition. Audio quality plateaus early; ladders are 2-4 rungs, never six.