

The four families, the 2026 facts, and the consent + copyright gates - on one page.

## Step 1 - which family is the feature?

**Neural codecs** Compress audio with a learned model. Ships in Lyra (Google Meet). Low legal risk; the cost is CPU / battery, not bandwidth.

**Neural cleanup** Noise suppression, echo cancel, packet-loss concealment. Repairs real audio. Ship today - but test what it deletes.

**Voice cloning / dubbing** Re-voice content in a known voice. ElevenLabs, Resemble. Works; the gate is consent + likeness rights, not engineering.

**Generative music / SFX** Create new audio from a prompt. Suno, Udio. The gate is training-data copyright - check the commercial terms.

## Step 2 - where it ships, what to watch

Family	Ships in 2026	Maturity	Main thing to watch
<b>Neural codecs</b>	Lyra in Google Meet	Speech prod.	CPU / battery cost
<b>Neural cleanup</b>	Every conferencing app	Mature	Erases real information
<b>Cloning / dubbing</b>	ElevenLabs, Resemble	Production	Consent; EU AI Act label
<b>Generative music</b>	Suno, Udio	Unsettled	Training-data copyright

## The 2026 facts that decide it

- One technique under all four: a neural network trained on huge audio. Existing sound = low risk; created / impersonated sound = where the law lives.
- Neural codecs cut bitrate 3-4x (SoundStream 3 kbps ~ Opus 12 kbps) but cost far more processor than Opus - bandwidth is nearly free, compute is not.
- AI dubbing (ElevenLabs v2) clones each speaker and conditions on their performance across 90+ languages; an hour of source dubs for low double-digit dollars.
- EU AI Act Art. 50 applies 2 Aug 2026: providers machine-mark synthetic audio; deployers disclose deepfakes. The date is fixed by Art. 113.
- Suno/Udio: Warner-Suno and Universal-Udio settled into licensing; Sony unsettled; a fair-use ruling is expected mid-2026. Generated-track status depends on the tool's terms.

## Before you ship - the gates

- Codec / cleanup: does the quality gain beat the per-client CPU and battery cost?
- Cleanup in clinical / security / music: test what the model removes, and give users an off switch.
- Cloning a real voice: get explicit, written, SCOPED consent (a course is not an advert).
- EU audience: label AI-generated audio and ensure machine-readable marking (Art. 50, from 2 Aug 2026).
- Generative music: use a generator with licensed / owned training data; read commercial terms; keep records.
- Any synthetic voice or audio: disclose it to the people who hear it.