

# Stable Audio 3.0: Open-Weight Music Generation Trained on Fully Licensed Data

Kabui, Charles

2026-05-23

---

[Read at ToKnow.ai](#)

---

The graphic features a dark blue background with a grid pattern. At the top left, the title 'Stable Audio 3.0: Open-Weight Music' is displayed in white and light blue. Below the title, a subtitle reads 'Four models trained on fully licensed data, three with open weights'. Three key metrics are highlighted in separate boxes: '6:20' (Max track length Medium model generation), '0.6B' (Parameters in Small model Runs on consumer hardware), and '<2s' (Generation time on H200 Seconds on MacBook Pro M4). The date 'May 23, 2026' is at the bottom left, and the 'ToKnow.ai' logo is at the bottom right.

## Stable Audio 3.0: Open-Weight Music

Four models trained on fully licensed data, three with open weights

- 6:20**  
Max track length  
Medium model generation
- 0.6B**  
Parameters in Small model  
Runs on consumer hardware
- <2s**  
Generation time on H200  
Seconds on MacBook Pro M4

May 23, 2026

ToKnow.ai

Stability AI released [Stable Audio 3.0](#), a family of four music and sound generation models trained entirely on licensed data from partnerships with [Universal Music Group](#) and [Warner Music Group](#). Three are open-weight on [Hugging Face](#): Small SFX and Small (both 0.6B parameters) for sound effects and full music composition on consumer hardware, and Medium (2B parameters) for longer, more complex tracks up to 6 minutes 20 seconds. The architecture

runs on a novel [semantic-acoustic autoencoder](#) (SAME) that compresses audio into a compact latent space, cutting generation to under two seconds on an H200 GPU and a few seconds on a MacBook Pro M4. Small generates up to two minutes of music, a 10x jump from 11 seconds in the previous version. LoRA fine-tuning (for customizing the model on your own audio) and audio inpainting (editing specific sections mid-track) are both supported.

Under the [Stability AI Community License](#), users own their outputs and can commercialize them freely (organizations above \$1M revenue need an Enterprise license). Most open music models either restrict commercial use or train on unlicensed data, which creates legal risk. A podcaster or indie game developer can download the 0.6B model, fine-tune it on their own audio library, and generate royalty-free music offline on a laptop. No API costs, no copyright uncertainty.

Music AI is splitting into two camps: proprietary platforms like [ElevenMusic](#) and [Google's Lyria 3](#) that control the pipeline end-to-end, and open-weight models you can run anywhere. Stable Audio 3.0 is the strongest entry in the open camp so far, and it comes with clear legal footing.

Sources:

- [Meet Stable Audio 3.0 \(Stability AI\)](#)
- [Stable Audio 3 Collection \(Hugging Face\)](#)
- [Stable Audio 3 Research Paper](#)
- [SAME Autoencoder Paper](#)
- [Stable Audio 3.0 Day-0 Support \(ComfyUI Blog\)](#)

---

*Disclaimer: For information only. Accuracy or completeness not guaranteed. Illegal use prohibited. Not professional advice or solicitation. Read more: [/terms-of-service](#)*