

The VibeVoice Setup Playbook

Clone Your Voice and Generate AI Audio for Free with Microsoft
VibeVoice

01

Get VibeVoice Running Locally

VibeVoice runs entirely on your own machine, so there's no subscription and no per-minute API fee once it's installed. Have Claude Code handle pulling the model and setting up the local environment — this is the step that normally trips up non-developers.

QUICK TIPS

- Ask Claude Code to install dependencies and start the local server for you
- Confirm the model loads and generates a test clip before moving on

02

Clone Your Voice From a 30-Second Clip

You only need about 30 seconds of clean, single-speaker audio to create a usable voice profile. Record it on your phone in a quiet room — no studio or special mic required.

QUICK TIPS

- Keep background noise and music out of the sample
- Speak at your normal pace and tone, not a performance voice

03

Generate Multi-Speaker Podcasts and Voiceovers

VibeVoice can produce full multi-speaker conversations with distinct custom voices in a single output, including in multiple languages. Feed it a script instead of NotebookLM and generate the whole episode locally.

QUICK TIPS

- Assign each speaker a separate cloned voice profile
- Generate in short sections first to check pacing before rendering a full episode

04

Automate Client Onboarding Audio

A cloned voice means every new client can get a personal-sounding welcome message without you recording it live. Write the script once, generate the audio automatically, and drop it into your onboarding sequence.

💡 QUICK TIPS

- Keep one 'template' script and swap in the client's name per generation
- Store generated clips with clear filenames so they're easy to reuse

05

Turn Blog Posts Into Audio

Paste any written post into VibeVoice and get a narrated version in your own voice, ready to publish alongside the text. This works well for repurposing content you've already written into a new format with almost no extra effort.

💡 QUICK TIPS

- Break long posts into shorter sections for cleaner pacing
- Re-listen to the first output before batch-generating older posts

Ready to go deeper?

This playbook is the short version. Inside the hub you get the full build walkthroughs, the prompt library, and the systems I use to ship this in a weekend.

Join the Build →

<https://hub.digicuratoragency.com/join>