



FrostyStack

Cool software that heats up your business

COOL SOFTWARE. HOT RESULTS.

A PRACTICAL GUIDE

Getting Started with ElevenLabs

The best AI voices, plus music, dubbing, video,
and voice agents. Updated for Eleven v4.

Text to Speech

Voice cloning

Eleven v4

Agents

Try ElevenLabs free >

Free plan with 10,000 credits a month. Paid plans from \$6/month.
Covers Eleven v4, audio tags, cloning, Flows, and the Claude MCP.



Affiliate disclosure: links in this guide are affiliate links. FrostyStack may earn a commission if you sign up, at no extra cost to you.
Independent guide by FrostyStack; not created or endorsed by ElevenLabs. Details verified October 2026.

Contents

What ElevenLabs Is in 2026	3
Plans and Credits	4
Choosing a Voice Model	5
Creating Your First Voiceover	6
Directing the Performance with Eleven v4	7
Voices: Library, Design, and Cloning	8
Studio: Audiobooks, Podcasts, and Long-Form Audio	9
Beyond Voice: Everything Else in ElevenCreative	9
ElevenAgents: Voice Agents for Your Business	10
Using ElevenLabs from Claude (MCP)	11
Rights, Safety, and Disclosure	12
Saving Credits and Avoiding Common Mistakes	12
Quick-Reference Checklist	13
Frequently Asked Questions	14

NOTE: How to use this guide

Most people come to ElevenLabs for its voices, so Sections 1 to 7 focus there: plans, choosing a model, your first voiceover, directing Eleven v4 with audio tags, cloning, and long-form audio in Studio.

Sections 8 to 10 cover everything else ElevenLabs now does: music, dubbing, image and video, Flows, voice agents, and running it from Claude. Current as of October 2026, including the Eleven v4 launch on September 28, 2026.

Affiliate disclosure: this guide is published by FrostyStack, an independent affiliate, and contains affiliate links. FrostyStack may earn a commission at no extra cost to you.

SECTION 1

What ElevenLabs Is in 2026

ElevenLabs made its name with the most realistic AI voices available. It has since grown into a full AI audio and media platform, but voice is still the core of everything it does.

ElevenLabs builds its own AI models for speech, transcription, music, and voice cloning, and now packages them into three products. Each runs on the same voice research, so a voice you clone or design works everywhere.

Product	What it is for	Highlights
ElevenCreative	Making content: voiceovers, audiobooks, podcasts, ads, dubbing, music, and video.	Text to Speech, Studio editor, Voice Library (17,500+ voices), cloning, Music, Sound Effects, Dubbing, Image and Video, and Flows
ElevenAgents	Conversational AI agents that talk, type, and take action.	Voice and chat agents across phone, web chat, email, and WhatsApp, with testing, guardrails, and analytics
ElevenAPI	Developers building ElevenLabs into their own apps.	Text to Speech, Speech to Text, Music, Dubbing, and Agents APIs and SDKs

What changed recently

When	Update
Sep 2026	Eleven v4 and Eleven v4 Turbo : ElevenLabs' most emotive voice models, in 90+ languages (up from about 70), with stronger audio-tag control and cloning.
May 2026	Dubbing v2 carries the original speaker's emotion across languages; Music v2 improves vocals and arrangement.
Mar 2026	Flows : a node-based canvas that chains image, video, voice, lip-sync, music, and sound effects into one repeatable pipeline.
Feb 2026	Expressive Mode for voice agents.
Jan 2026	Scribe v2 , ElevenLabs' most accurate transcription model.

Hear the difference yourself

The free plan includes 10,000 credits a month, enough to try the voices and Eleven v4 before you pay.

[Try ElevenLabs >](#)

SECTION 2

Plans and Credits

ElevenLabs runs on monthly credits. For standard speech, one credit is roughly one character of text, so 1,000 credits is about a minute of audio.

Plan	Monthly	Credits/month	Key unlocks
Free	\$0	10,000	Try every voice; Voice Design. No commercial use; you must credit ElevenLabs
Starter	\$6	30,000	Commercial license, Instant Voice Cloning
Creator	\$22	121,000	Professional Voice Cloning (1 voice), higher-quality audio, usage-based extra credits
Pro	\$99	600,000	High-fidelity 44.1 kHz audio via API, higher throughput
Scale	\$299	1.8 million	3 seats for small teams
Business	\$990	6 million	10 seats, lowest per-minute rates

Prices as of October 2026; Enterprise is custom. Annual billing gives about two months free (for example, Creator drops to about \$18.33 a month). ElevenLabs often runs first-month promotions, such as a discounted first month on Starter or Creator; see [current ElevenLabs pricing](#) for today's offers.

Which plan should you pick?

If you...	Choose
Just want to test voices	Free
Will publish or monetize anything, even occasionally	Starter: you are mainly buying the commercial license
Make regular content or want a high-fidelity clone of your own voice	Creator: about two hours of audio a month
Produce audiobooks, daily content, or need uncompressed audio	Pro
Share an account across a team	Scale or Business

WATCH OUT: Free-plan audio is not for business use

On the Free plan you cannot use the audio commercially and must attribute ElevenLabs when you publish. If a voiceover goes into an ad, a client project, or a monetized video, you need at least Starter.

Voice agents are metered separately in call minutes (for example, 15 minutes on Free and 275 on Creator), so heavy agent use draws on a different allowance than voiceovers.

SECTION 3

Choosing a Voice Model

ElevenLabs offers several speech models. Picking the right one is the single easiest way to get better audio.

Model	Best for	Notes
Eleven v4	Voiceovers, ads, audiobooks, characters, multi-speaker dialogue	Most emotive and highest quality; 90+ languages; follows audio tags closely; up to 10,000 characters per generation
Eleven v4 Turbo	Voice agents and fast drafts	Same expressive style at about 100 ms latency
Eleven Multilingual v2	Long narration that must stay very consistent	ElevenLabs' long-standing lifelike, stable model
Eleven Flash v2.5	Real-time apps where speed matters most	About 75 ms latency
Eleven v3	Older projects	Introduced audio tags; v4 now does them better

TIP: Start with Eleven v4

For almost any creative work, start with Eleven v4. Switch to Multilingual v2 only if you need a long, very even read, and to v4 Turbo or Flash only for live, low-latency uses such as agents.

SECTION 4

Creating Your First Voiceover

Text to Speech is the heart of ElevenLabs. You can have a finished voiceover in a few minutes.

- 1 **Sign up** at elevenlabs.io and open **Text to Speech** in ElevenCreative.
- 2 **Pick a voice.** Browse the **Voice Library** by use case (narration, advertisement, conversational, characters, social media), language, accent, and age. Play samples and add favorites to My Voices.
- 3 **Choose the model:** Eleven v4 for most work (Section 3).
- 4 **Paste your script.** Write the way people speak: short sentences, contractions, and numbers written the way you want them read.
- 5 **Adjust settings if needed.** **Stability** trades consistency (higher) against expressiveness (lower); **Similarity** controls how closely the output sticks to the original voice.
- 6 **Generate, listen, and regenerate** any line that sounds off. Each generation uses credits, so fix the script first rather than re-rolling repeatedly.
- 7 **Download** the audio as MP3 or WAV and drop it into your video or podcast editor.

TIP: Spend credits on the final read, not the experiments

Test a voice and style on a short paragraph first. Once it sounds right, generate the full script. For long pieces, work in Studio (Section 7) so you can regenerate one sentence instead of the whole thing.

Pronunciation fixes

If a name or brand is mispronounced, try spelling it phonetically in the script ("FrostyStack" might become "Frosty Stack"). Eleven v4 also supports IPA phonetic spelling for precise control, and pronunciation dictionaries let you save fixes for reuse.

SECTION 5

Directing the Performance with Eleven v4

Older text-to-speech models read text aloud. Eleven v4 performs it. You direct the performance two ways: with inline audio tags and with plain-language direction.

Audio tags

Put short instructions in square brackets right where you want them. The voice acts them out instead of reading them. Tags can control emotion, delivery, and even background sounds:

```
[whispers] Don't look now... [pause] but someone's been reading your emails.  
[laughs] Okay, okay, I'll admit it. [sighs] That was my fault.  
[said angrily in a French accent] I asked for the window seat!  
[light rain] [phone buzzing] She ignored it and kept walking.
```

Tips for tags that work

- **Match the tag to the voice.** A calm narrator voice will not shout convincingly; pick a voice whose range fits.
- **Use tags sparingly.** One or two per paragraph sounds natural; one per sentence sounds like an act.
- **Let the text carry emotion too.** v4 reads tone and context from the words, so a well-written line needs fewer tags.
- **Use punctuation for pacing.** Ellipses and commas slow the read. Note that Eleven v4 has no speed slider and does not read SSML break tags, so pace with the writing itself.

Multi-speaker dialogue

Eleven v4 is built to handle natural conversation between several speakers, with consistent voices across long productions. Use it for podcast-style dialogue, interviews, character scenes, and two-voice ads.

TIP: Same voice, other languages

Eleven v4 can switch languages and accents while keeping the rest of a voice's character. That makes it practical to keep one brand voice across markets.

SECTION 6

Voices: Library, Design, and Cloning

You have four ways to get a voice, from picking one off the shelf to cloning your own.

Option	What it is	Plan
Voice Library	17,500+ voices shared by ElevenLabs and its community, searchable by use case, language, and accent	All plans
Voice Design	Describe a voice in words ("warm, gravelly older man, slight Southern accent") and ElevenLabs creates it	All plans, including Free
Instant Voice Clone	A quick clone from a short sample; on Eleven v4, about 10 seconds of clean audio	Starter and up
Professional Voice Clone	A high-fidelity clone trained on a longer set of recordings; the closest match to the real voice	Creator and up

Getting a great clone

- Record in a quiet room with a decent microphone; no music, echo, or background noise.
- Speak the way you want the clone to sound. A flat read produces a flat clone.
- For a Professional clone, provide more audio, ideally well over 30 minutes, in a consistent style.
- Use the **Voice Isolator** to clean up a noisy recording before cloning it.

WATCH OUT: Clone only voices you have the right to use

Cloning requires the speaker's consent. Clone your own voice, or a voice actor's with written permission. ElevenLabs moderates content, verifies Professional clones, and adds provenance signals so AI audio can be identified. Misuse can get your account banned and may break the law.

Voice actors can also share Professional clones in the Voice Library and earn when others use them, and the Iconic Marketplace offers licensed voices of well-known figures.

SECTION 7

Studio: Audiobooks, Podcasts, and Long-Form Audio

Studio is ElevenLabs' all-in-one editor for long projects. Import a manuscript, script, or article, and it splits the text into chapters and paragraphs you can voice, edit, and regenerate one piece at a time.

- **Assign voices** to narrators and characters, then switch voices per paragraph or line.
- **Regenerate single sentences** without redoing a whole chapter, which saves credits.
- **Add music and sound effects** on a timeline, and export finished audio or video.
- **Podcasts:** turn a document or URL into a two-host conversation, then edit the script before generating.

TIP: Audiobook workflow

Import the manuscript, choose a narrator voice with Eleven v4, generate one chapter, and listen carefully on headphones. Fix pronunciations and pacing in the text, then generate the rest. Check your distributor's rules on AI-narrated books before publishing.

SECTION 8

Beyond Voice: Everything Else in ElevenCreative

ElevenLabs is no longer just a voice tool. These features live alongside Text to Speech in the same account and use the same credits.

Tool	What it does
Music	Studio-quality songs or instrumentals from a text prompt, in any genre. Built with artists, labels, and publishers and cleared for commercial use, with rights that vary by plan.
Sound Effects	Custom effects, ambience, and soundscapes from a description, plus a searchable SFX library.
Dubbing	Translate a video into other languages in the original speaker's voice. Dubbing v2 carries the speaker's emotion and performance across languages.
Speech to Text	Scribe v2 transcription with speaker labels and timestamps, for captions, notes, and show notes.
Voice Changer	Re-perform your own recording in a different voice, keeping your timing and emotion.
Voice Isolator	Remove background noise and music from a recording.

Tool	What it does
Image and Video	Generate images and turn ideas into videos with leading models such as Veo, Kling, Wan, and Seedance.
Flows	A node-based canvas that chains image, video, Text to Speech, lip-sync, music, and sound effects into one reusable pipeline, with 35+ image and video models.

A Flows example: one product, many ads

- 1 In ElevenCreative, open **Flows** and click **New Flow**. Name it as a reusable template.
- 2 Add an **Image Generation** node with photos of your product as references.
- 3 Connect a **Video Generation** node and describe the motion ("slow push-in, soft light").
- 4 Add **Text to Speech** for a voiceover and **Music** for a bed, then combine them.
- 5 Swap the product, hook, voice, or language to generate variations without rebuilding.

Video generation in Flows requires a paid plan. If video is your main focus, compare with dedicated video platforms such as [Higgsfield](#), which FrostyStack also covers.

SECTION 9

ElevenAgents: Voice Agents for Your Business

ElevenAgents lets you build AI agents that hold natural conversations by voice or text, in 90+ languages, across phone, website chat, email, and WhatsApp. Typical uses are answering customer questions, booking appointments, qualifying leads, and handling after-hours calls.

Building a simple agent

- 1 Open **Agents** and create a new agent from a template or from scratch.
- 2 Write its instructions: who it is, what it should do, what it must never do, and when to hand off to a human.
- 3 Choose its voice (v4 Turbo or another low-latency model for natural back-and-forth) and add a knowledge base, such as your FAQ, services, and hours.
- 4 Connect tools if needed, such as a calendar or CRM, and set **guardrails** for compliance.
- 5 Use **Testing** to simulate conversations before going live, then connect a phone number or embed the chat widget.
- 6 Watch **Analytics** (resolution rate and other metrics) and refine the instructions over time.

WATCH OUT: Agents need careful setup

An agent speaks for your business. Test it thoroughly, tell callers they are talking to AI where the law requires it, and get consent before recording calls where required.

Agent call minutes are a separate allowance on each plan (Section 2), and live calls also depend on concurrency limits.

SECTION 10

Using ElevenLabs from Claude (MCP)

ElevenLabs maintains an official MCP server, so Claude can generate speech, clone voices, transcribe audio, and more from a chat.

MCP (Model Context Protocol) is an open standard that lets an AI assistant use outside tools. With the ElevenLabs server connected, you can ask Claude to write a script *and* voice it in one conversation.

Set it up in Claude Desktop

- 1 In ElevenLabs, create an **API key** under Developers > API Keys.
- 2 Install **uv**, a small Python tool runner that the server uses (instructions at astral.sh/uv).
- 3 In Claude Desktop, open **Settings > Developer > Edit Config** and add the server below, with your key.
- 4 Restart Claude Desktop. ElevenLabs tools appear in the chat's tool list.

```
{ "mcpServers": {  
  "ElevenLabs": {  
    "command": "uvx",  
    "args": ["elevenlabs-mcp"],  
    "env": { "ELEVENLABS_API_KEY": "your-api-key-here" }  
  }  
}
```

This server runs on your own computer, so it works in the Claude Desktop app rather than claude.ai in a browser. Generated audio files are saved to your computer. ElevenLabs entries for its Agents product have also started appearing in Claude's connector directory; check there for the latest one-click options.

Things to ask Claude

You say	What happens
"Write a 30-second ad for my coffee shop and voice it with a warm, friendly male voice."	Claude writes the script and generates the MP3.
"Transcribe this interview recording with speaker labels."	Claude runs Speech to Text and returns the transcript.
"Design a voice for a wise old wizard and read this paragraph."	Claude uses Voice Design, then Text to Speech.
"Create a 10-second sound effect of waves on a rocky beach."	Claude generates the effect.

WATCH OUT: Claude spends your credits

Everything Claude generates uses your ElevenLabs credits at normal rates. Keep your API key private, and ask Claude to confirm before long or batch generations.

SECTION 11

Rights, Safety, and Disclosure

- **Commercial use starts on Starter.** Free-plan audio cannot be used commercially and must credit ElevenLabs.
- **Consent for every clone.** Never clone a real person's voice without permission, including celebrities, politicians, or clients' staff.
- **Disclose AI where it matters:** on platforms that require AI labels, in political or news content, and to callers speaking with an AI agent.
- **Music rights vary by plan.** Check the license tier before using generated music in ads or client work.
- **Platform rules apply.** YouTube, TikTok, audiobook distributors, and ad networks each have their own policies on synthetic media.

NOTE: ElevenLabs' safety measures

ElevenLabs moderates generated content, enforces consequences for misuse, and supports provenance so audio can be identified as AI-generated. Following its rules protects your account as well as other people.

SECTION 12

Saving Credits and Avoiding Common Mistakes

Mistake	Better approach
Regenerating a whole script to fix one line	Use Studio and regenerate only the sentence that needs it.
Testing voices on long scripts	Test on one short paragraph, then generate the full piece.
Publishing Free-plan audio commercially	Upgrade to Starter before anything goes public.
Overusing audio tags	One or two per paragraph; let the writing carry the emotion.
Cloning from noisy recordings	Record cleanly or run the Voice Isolator first.
Using v4 for a live agent	Use v4 Turbo or Flash for real-time conversations.
Ignoring pronunciation	Fix names with phonetic spelling, IPA, or a pronunciation dictionary.

SECTION 13

Quick-Reference Checklist

Print this page and keep it handy.

Getting started

- ☐ Account created via [the FrostyStack ElevenLabs link](#); plan chosen (Starter or higher for commercial use)
- ☐ Favorite voices saved from the Voice Library, or a voice designed or cloned (with consent)
- ☐ Default model set to Eleven v4 for creative work

Every voiceover

- ☐ Script written for the ear: short sentences, numbers spelled out as spoken
- ☐ Tested on a short paragraph; stability and similarity adjusted
- ☐ Audio tags added sparingly; pronunciations fixed
- ☐ Listened on headphones before downloading

Before publishing

- ☐ Commercial license confirmed for your plan
- ☐ Consent on file for any cloned voice
- ☐ AI disclosure added where the platform or law requires it

SECTION 14

Frequently Asked Questions

Is ElevenLabs free?

Yes, there is a free plan with 10,000 credits a month, about 10 minutes of speech. Free audio cannot be used commercially. [Try it free.](#)

What is the best model?

Eleven v4, launched September 28, 2026, is ElevenLabs' most emotive model and ranked first on the Artificial Analysis voice leaderboard at launch. Use v4 Turbo for real-time agents.

How many languages are supported?

Eleven v4 speaks 90+ languages.

Can I clone my own voice?

Yes. Instant Voice Cloning starts on Starter; Professional Voice Cloning starts on Creator.

Do unused credits roll over?

Not on the Free plan. Paid-plan rollover rules vary, so check your plan details before counting on it.

Can I make music and sound effects?

Yes. Music and Sound Effects are part of ElevenCreative; music commercial rights vary by plan.

Can ElevenLabs answer my business phone?

Yes, with ElevenAgents. Agent minutes are a separate allowance on each plan.

Does it work with Claude?

Yes, through ElevenLabs' official MCP server in Claude Desktop (Section 10).

Give your content a voice with ElevenLabs

Start free with 10,000 credits a month, then upgrade when you are ready to publish.

[Try ElevenLabs free >](#)



FrostyStack

Cool software that heats up your business.

FrostyStack reviews practical AI tools, automation, and growth strategies that help businesses save time, increase traffic, and grow faster. Find more tool guides and growth playbooks at frostystack.com.

Affiliate disclosure: links to ElevenLabs and Higgsfield in this guide are affiliate links. If you sign up through them, FrostyStack may earn a commission at no extra cost to you. FrostyStack is an independent affiliate; this guide was not created or endorsed by ElevenLabs. Details verified October 2026 from ElevenLabs' website, documentation, and published pricing; features and prices change often. Nothing in this guide is legal advice.