




SYNTHESYS

Studio Complete Training

Every model, direct. Make images, video clips and voiceovers in one place — pick the model, see the price before you press, and pay only for what arrives.

SEPTEMBER 2026 · CUSTOMER EDITION · EVERY SCREEN IS A
CAPTURE OF THE REAL APP



PART 1 · START HERE

- 1.1 Meet Studio
- 1.2 The Studio screen
- 1.3 How credits work in Studio
- 1.4 History & How it works

PART 2 · IMAGE

- 2.1 Your first image
- 2.2 Choose a model
- 2.3 Image models at a glance
- 2.4 Size, quality & how many
- 2.5 References — lock a product or a face
- 2.6 @ mentions
- 2.7 Looks
- 2.8 Edit an image
- 2.9 Your images

PART 3 · VIDEO

- 3.1 Your first clip
- 3.2 Choose a video model
- 3.3 Video models at a glance
- 3.4 What a clip costs
- 3.5 Start & end frames
- 3.6 Length, resolution, aspect & audio
- 3.7 Video Looks
- 3.8 References & reference audio
- 3.9 Rendering & your clips

PART 4 · VOICEOVER

- 4.1 Your first voiceover
- 4.2 Write & price the script
- 4.3 Pick a voice
- 4.4 Voiceover models at a glance
- 4.5 Clone your voice
- 4.6 Design a voice
- 4.7 Your voiceovers

PARTS 5–6

- 5 Tips & tricks
- 6 FAQ & troubleshooting

PART 1

Start here

Studio is where you make single pieces — one image, one clip, one voiceover — with the model of your choice. Four lessons: what it is, where everything sits, how credits work, and where your results go.


Meet Studio

The Video Agent builds whole videos for you. **Studio** is the opposite: you pick the exact model, set every option yourself, and get one result back — fast.

- 1 **Image** — 18 image models from 12 makers, from 1 credit a picture up to premium 4K.
- 2 **Video** — 10 video models: text-to-video, animate a still, clips up to 30 seconds.
- 3 **Voiceover** — 7 voice models, hundreds of voices, plus your own cloned or designed voice.
- 4 **Open it** from the sidebar: **Studio** (marked NEW), right under New chat.

TIP

Everything you make in Studio is saved to **Files** automatically — and any image or clip can go straight into a Video Agent chat with **Use in a video**.




The Image tool. Controls on the left, your History on the right.

The Studio screen

Every tool uses the same layout, so once you know one you know all three.

- 1 **The tool tabs** — Image · Video · Voiceover, at the top of the left panel. Each tab keeps its own prompt.
- 2 **Inputs** — references, frames or a voice. Studio only shows the slots the chosen model can actually use.
- 3 **The prompt** — Prompt, Shot or Script, depending on the tool. ⌘/Ctrl + Enter generates.
- 4 **The bottom bar** — size and length settings, the **model chip** with its live price, and **Generate**.
- 5 **The right pane** — **History** and **How it works**, with your credit balance in the top corner.



The Video tool — same layout as Image. Change a model and the inputs change with it — a model that takes no images simply hides the image slots.

How credits work in Studio

You always see the price before you press anything, and you only pay for results that actually arrive.


- 1 Priced everywhere** — on the model chip, on every row of the model list, and on every size or quality option.
- 2 Held, then charged** — pressing Generate holds the credits; they come off only when the result lands in your History.
- 3 Failed runs are free** — every failure card says *No credits charged*. In a batch, only the images that arrive are charged.
- 4 Short on credits?** Nothing starts. Studio shows how many you're short, the models that fit your balance, and **Buy credits**.

CREDITS

Prices depend on the model and its settings. Examples: a FLUX.1 [schnell] image from **1 credit**, a 5-second MiniMax clip from **25**, a voiceover from **1**.

TIP

Lost the connection mid-run? Check History before pressing Generate again — if the result still arrives, it lands there and is charged once.



Every option shows its own price — here, GPT Image 2 at Low, Medium and High.


History & How it works

The right pane is your workbench. **History** keeps everything you've made with the current tool, newest first.

- 1 List or Grid** — switch views at the top right; the slider beside it makes the previews bigger or smaller.
- 2 Each result** shows the model, your prompt, the inputs you used, the settings and what it cost.
- 3 Reuse settings** puts that run back into the form — model, prompt, references and all. It never runs by itself.
- 4 How it works** — a three-step picture guide for each tool, always one click away.

TIP

Grid view is the fastest way to scan a big session. Click any tile to jump to it in the list.



Grid view of the Image History.

PART 2

Image

Product shots, portraits, illustrations and edits — with 18 models to choose from. Nine lessons, from your first picture to editing one you already have.

Your first image

- 1 Open **Studio** and stay on the **Image** tab (Create is selected).
- 2 **Describe the image** — the subject, the light, the lens. “A clay fox reading under a mushroom lamp, soft rim light, 50mm.”
- 3 **Pick a model and a size** — the default is GPT Image 2 at Medium, square: **13 credits**.
- 4 Press **Generate**. Most images arrive in well under a minute.

TIP

Stuck for words? Click a **Look** under the prompt — with an empty prompt it writes a full description for you to edit (lesson 2.7).

History

How it works

Make an image in three steps

Every model, direct — and you see the credits before you press anything.



+

A clay fox reading under a mushroom lamp, soft rim light, 50mm

 **GPT Image 2**
Low to High · 3 refs3–44

 **Nano Banana 2**
up to 4K · 14 refs**14–31** ✓

 **Ideogram v4**
Medium · text only7



STEP 1

Describe the image

The subject, the light, the lens. Add references to lock a product or a face — up to 14 on some models.

STEP 2


Pick a model and a size

Filter by maker, or let the tiers guide you. The price is on every row.

STEP 3

Get your image

Usually under a minute. Download it, or use it as the first frame of a video.




Don't know where to start?
Open Training — short lessons on prompts, looks and models.

>

Choose a model

Click the **model chip** next to Generate. The list groups models into three tiers and shows each one's price.

- 1 **Tiers** — **Fast & cheap**, **Balanced** and **Highest quality**.
- 2 **Filter by maker** — Featured, All 18, then Google, OpenAI, ByteDance, Black Forest Labs and more.
- 3 **Read the row** — “up to 4K · 14 refs” tells you the top size and how many reference images it takes; “text only” means no images in.
- 4 **Prices** — a range like 14–31 covers the model's size or quality options. The cheapest model wears a **CHEAPEST** badge.
- 5 **Search** by model name, maker or ability — try “editing” or “references”.



The full list, grouped by tier, with the price range on the right of every row.

Image models at a glance

MODEL	MAKER	WHAT IT OFFERS	CREDITS / IMAGE	REFS	EDIT
FLUX.1 [schnell]	Black Forest Labs	Fast & cheap · up to 2K · text only	1 – 3	—	—
MiniMax Image-01	MiniMax	Fast & cheap · text only	2	—	—
FLUX.2 [dev]	Black Forest Labs	Fast & cheap · up to 2K · text only	3 – 10	—	—
Grok Imagine Image	xAI	Fast & cheap · up to 2K	4 – 6	3	✓
Ideogram v4	Ideogram	Fast & cheap · text inside the image	7	—	—
Kling Omni 3	Kuaishou	Balanced · up to 4K	6 – 12	10	✓
FLUX.2 [pro]	Black Forest Labs	Balanced · up to 2K · text only	6 – 15	—	—
Recraft V4.1	Recraft AI	Balanced · text only	7	—	—
GPT Image 2.5 Flare	OpenAI	Balanced · Low to High · 1:1, 16:9, 4:3 only	4 – 24	—	—
Seedream 4.5	ByteDance	Balanced · 2K	8	10	✓
FLUX.1 Kontext [pro]	Black Forest Labs	Balanced · edits from one reference	8	1	✓
Bria FIBO Gen 1.5	Bria AI	Balanced · up to 4MP	8	4	✓
Nano Banana 2	Google	Balanced · up to 4K	14 – 31	14	✓
Qwen Image 3	Alibaba	Balanced · up to 2K	15	3	✓
GPT Image 2 ★	OpenAI	Highest quality · Low to High · the default	3 – 44	3	✓
Seedream 5.0 Pro	ByteDance	Highest quality · up to 2K	14 – 36	10	✓
Hunyuan Image 3.0 Instruct	Tencent	Highest quality · up to 2K	18 – 72	3	✓
Nano Banana Pro	Google	Highest quality · 1K or 4K	30 – 60	14	✓

Refs = the most reference images the model takes. Edit = available in Edit mode. Ranges cover size, quality and — for Seedream 5.0 Pro and Grok Imagine Image only — the number of references. The exact price is always on the model chip before you press Generate.



Size, quality & how many

Three settings sit in the bottom bar, left of the model chip.

- 1 **Aspect ratio** — eight shapes from 16:9 Widescreen to 9:21 Ultra-tall. Each model lists only the shapes it makes.
- 2 **Resolution or Quality** — 1K / 2K / 4K, or Low / Medium / High. Every option shows its own price.
- 3 **Images** — ×1 to ×4 per press. Four images cost four times one image; each is made separately, all at once.
- 4 Switch to a model that can't do your setting and Studio picks the nearest one and tells you what changed.

CREDITS

On GPT Image 2, GPT Image 2.5 Flare and Seedream 5.0 Pro the **shape changes the price** — e.g. Seedream 5.0 Pro at 2K costs 14 in 16:9 but 27 in 1:1.



Left: the eight aspect ratios. Right: 1 to 4 images per press.


References — lock a product or a face

Give the model real pictures to follow — your product, a face, a colour, a style. How many depends on the model: 1, 3, 4, 10 or up to **14**.

- 1** Click **+** in References to open your library: **Uploads**, **Image generations** (everything you made in Studio) and **Stock photos**.
- 2** Pick **several at once** — tiles get a number in the order you chose them — then **Use N images**.
- 3** Or **drop and paste** — drag files anywhere onto the form, or paste an image straight from your clipboard.
- 4** Uploads are saved to your library, so you can reuse them next time.

TIP

Switching to a model with fewer slots doesn't delete anything — the extra images are **held** and come back when you switch to a model that takes them.



The library window, with two past generations picked as references for Nano Banana 2.


@ mentions — point at a reference

With two or more references, tell the model which one is which. Type @ in the prompt.

- 1 Type @ at the start of the prompt or after a space — a list of your attached references opens.
- 2 Use ↑ / ↓ and **Enter** (or click) to drop a chip into your sentence.
- 3 Write around the chips: “A frosted amber bottle like **[bottle]**, standing on moss with the texture of **[leaf]**.”
- 4 Press **Backspace** once to select a chip, twice to remove it. Removing a reference removes its chips too.

TIP

The grey text in the list — “reference image 1” — is exactly what the model reads, so you always know which picture you’re pointing at.



Two references attached; @ opens the list of what you can mention.


Looks — a style in one click

A Look is a ready-made art direction. There are **30** image Looks in four families.

- 1 Photography (13)** — Product shot, Portrait, Flat lay, Lifestyle, Food table, Golden hour, Cinematic still, Neon night...
- 2 3D & render (6)** — 3D product render, Clay render, Isometric scene, 3D character, Glass & chrome, Paper craft.
- 3 Illustration (6) and Graphic (5)** — watercolour, anime cel, comic halftone, flat vector, bold poster, retro print...
- 4 Empty prompt?** The Look writes the whole prompt for you to edit. **Already typed?** Your words stay and the Look's direction is added when you generate.

TIP

Click **Browse all 30** to search and filter them. Click the chosen Look again, or its X, to remove it.



The Looks browser — every tile is a real image made with that Look.

Edit an image

Switch **Create** to **Edit** to change a picture you already have instead of making a new one.

- 1 Image to edit** — drop, paste or pick the picture from your library. It's required.
- 2 Add images** (optional) — a logo, a product or a colour to bring into the picture.
- 3 Describe the change** — “make the jacket red”, “remove the person on the left”.
- 4** The shape **follows the image**. Only the 11 models that can edit are listed.

TIP


Every image in History has **Edit this image** — one click puts it in the Edit slot. Edited results show an *Edit* tag and the picture they came from.

Edit mode: the image to edit, optional extra images, and the change you want.

Your images

Every image lands in History and in **Files** — there's no save button to forget.

- 1 **Use in a video** — opens a new Video Agent chat with the image attached: “Make a video featuring this image.” Nothing runs until you send it.
- 2 **Edit this image** — straight into Edit mode (lesson 2.8).
- 3 **Reuse settings** — refill the form with this run, change one thing, and generate again.
- 4 **Download** — the full-size file. Click any picture to see it full screen.
- 5 **For a clip** — on the Video tab, pick it as a start frame from the library's **Image generations** tab.



A History row: model, prompt, inputs, settings, price — and what you can do next.

PART 3

Video

Ten video models — type a shot, or bring a still to life. Nine lessons: models, prices, frames, settings, Looks, references and what happens while a clip renders.

Your first clip

- 1 Open the **Video** tab.
- 2 **Add a start frame** (optional) — drop an image on the **First frame** tile or pick one from your library. The clip opens on it.
- 3 **Describe the shot** — the subject, the motion, the camera — or click a **Look** to start from a proven style.
- 4 Check **Resolution · Aspect ratio · Length** in the bottom bar, then press **Generate**.

CREDITS

The default model, **MiniMax H3 Max Turbo**, makes a 5-second 480p clip for **25 credits** — the cheapest way to test an idea.

TIP

Clips usually take **1–5 minutes**. You can close the tab — it keeps rendering and lands in History and Files.

History How it works

Make a clip in three steps

Every model, direct — and you see the credits before you press anything.

+

First frame




STEP 1
Add a start frame
Drop an image on the tile, or pick one from your library. Some models work from words alone.




Luxury dreamscape

STEP 2
Choose a look
A look fills in the style for you — or skip it and describe the shot yourself.
[Browse looks](#)



STEP 3
Get your clip
It lands in History and in Files. Credits come off only when it arrives.



Don't know where to start?
Open Training — short lessons on prompts, looks and models.


[>](#)

How it works for Video: add a start frame, choose a look, get your clip.

Choose a video model

Click the **model chip**. Rows are grouped **Fast & cheap**, **Balanced** and **Highest quality**; makers are Google, ByteDance, MiniMax, Black Forest Labs and xAI.

- 1 **Read the row** — “up to 1080p · 4–30s · 3 refs · start frame” tells you the top resolution, the lengths, how many reference images and whether it takes a start frame.
- 2 **The price** on each row is for a 5-second clip at 16:9 — it moves with length and resolution.
- 3 **“needs a start frame”** (Grok Imagine) means it only animates an image. **“text only”** (Veo 3 Fast) means words alone.
- 4 Switching models keeps what it can — anything the new model can’t use is held, and comes back when you switch again.



Featured video models: MiniMax H3 Max Turbo (the cheapest) and Seedance 2.5 (up to 30 seconds).

Video models at a glance

MODEL	MAKER	RESOLUTION	LENGTH	START / END FRAME	REFS	SOUND
MiniMax H3 Max Turbo ★	MiniMax	480p – 1080p	5 – 15s	✓ / ✓	—	—
Seedance 2.0 Mini	ByteDance	480p – 720p	4 – 15s	✓ / ✓	3	Audio switch
Grok Imagine	xAI	720p	6 – 15s	required / —	—	Always on
Seedance 2.0 Fast	ByteDance	480p – 720p	4 – 15s	✓ / ✓	3	Audio switch
MiniMax H3 Max	MiniMax	480p – 1080p	5 – 15s	✓ / ✓	—	—
Seedance 2.0	ByteDance	480p – 1080p	4 – 15s	✓ / ✓	3	Audio switch
Gemini Omni	Google	720p – 1080p	4 – 10s	— / —	3	Always on
Veo 3 Fast	Google	1080p · 16:9	4 – 8s	words only	—	Always on
FLUX.3	Black Forest Labs	720p – 1080p	5 – 15s	✓ / —	—	Audio switch
Seedance 2.5 ★	ByteDance	480p – 1080p	4 – 30s	✓ / ✓	3	Audio switch

In the order the model list shows them: Fast & cheap, then Balanced, then Highest quality. ★ = Featured. Seedance clips longer than 15s are 480p or 720p. 4:5 Social portrait is on the Seedance models only. Gemini Omni renders 16:9 and 9:16 natively; other shapes are reframed at 720p for +12 credits a second.

What a clip costs

Credits per clip for common choices. Turning **Audio** on or off doesn't change the price.

MODEL	CHEAPEST CLIP	10S AT 720P	10S AT 1080P	LONGEST CLIP
MiniMax H3 Max Turbo	25 · 5s 480p	80	160	240 · 15s 1080p
Seedance 2.0 Mini	40 · 4s 480p	220	—	330 · 15s 720p
Grok Imagine	60 · 6s 720p	100	—	150 · 15s 720p
Seedance 2.0 Fast	62 · 4s 480p	330	—	495 · 15s 720p
MiniMax H3 Max	50 · 5s 480p	160	320	480 · 15s 1080p
Seedance 2.0	76 · 4s 480p	410	1,020	1,530 · 15s 1080p
Gemini Omni	90 · 4s	180	180	180 · 10s
Veo 3 Fast	120 · 4s 1080p	—	—	240 · 8s 1080p
FLUX.3	170 · 5s 720p	340	580	870 · 15s 1080p
Seedance 2.5	112 · 4s 480p	630	1,580	1,890 · 30s 720p

Gemini Omni costs the same at 720p and 1080p in 16:9 or 9:16. The exact price for your settings is always on the model chip and on every Resolution option before you press Generate.

Start & end frames

Frames are the strongest control you have: the clip **opens** on your start frame and, where supported, **ends** on your end frame.

- 1 Add a frame** — drop or paste an image on the tile, or pick from your library. Your Studio images are under **Image generations**.
- 2 End frame** — on MiniMax and Seedance. On Seedance it needs a start frame too; MiniMax takes one on its own.
- 3 Grok Imagine** needs a start frame — it animates the image you give it.
- 4 MiniMax** follows your frame's shape — the aspect setting is ignored once a frame is set.

TIP

Make the perfect still on the Image tab first, then animate it here. It's the cheapest way to control exactly what the clip looks like.

START FRAME optional

×

First frame



END FRAME optional

+

Last frame

REFERENCES

+

Up to 3 — lock a look, a product or a face

REFERENCE AUDIO

+

 Add audio

Remove the start frame to use audio

SHOT

Slow push-in on the amber bottle, morning light sweeping across the backdrop, a soft reflection blooming on the glass

A Studio image as the start frame, with the shot described below it.



Length, resolution, aspect & audio

Four settings in the bottom bar. Each model only offers what it can make.

- 1 **Resolution** — 480p, 720p or 1080p. Each option shows its price for your length and shape.
- 2 **Aspect ratio** — up to eight shapes: 16:9, 9:16, 1:1, 4:5, 4:3, 3:4, 21:9, 9:21.
- 3 **Length** — 4 to 15 seconds on most models; **Seedance 2.5** goes to 20, 25 and 30.
- 4 **Audio on / off** — on Seedance and FLUX.3. Veo, Gemini Omni and Grok always make sound; MiniMax has no audio option.

TIP

Seedance models start at 4 seconds, most others at 5. If your length isn't available, Studio switches to the nearest one and says so under the bar.




Left: resolution with prices for a 5s clip. Right: the aspect menu on Seedance 2.5.

Video Looks

About **55** video Looks in three families — **Commercials**, **UGC** and **Graphics & painted**. Hover any Look to play its preview clip.

- 1 Click a quick pick under Shot, or **Browse all** to search and filter every Look.
- 2 **Empty Shot box?** The Look writes the shot for you — edit it to make it yours.
- 3 **Already wrote your shot?** Your words stay; the Look's direction is added when you generate.
- 4 Pair a Look with a start frame: the frame decides *what* is in the clip, the Look decides *how* it moves and feels.



The video Looks browser — Commercials, UGC and Graphics & painted.

References & reference audio

Seedance models and Gemini Omni take up to **3 reference images** to lock a product, a face or a look. Seedance can also follow **reference audio**.

- 1 References** — add them like on the Image tab, and use @ in the Shot box to point at one.
- 2 Reference audio** (Seedance) — **Add audio**, or **use one of your voiceovers** from the Voiceover tab.
- 3** Audio clips are WAV or MP3 up to 15 MB. Seedance 2.0 models take up to 3 clips (15s in total); Seedance 2.5 takes up to 10.
- 4** Some inputs don't mix: reference audio can't be used with a frame, and on Seedance an end frame can't be used with reference images.

TIP

Studio tells you what clashes as you go — “Remove the start frame to use audio” — so you never press Generate on a clip that can't run.

START FRAME optional

END FRAME optional

First frame

Last frame

Add a start frame first

REFERENCES

+

Up to 3 — lock a look, a product or a face

REFERENCE AUDIO

+ Add audio

use one of your voiceovers

Up to 10 — a voice or a sound for the clip to follow

SHOT

Describe your shot — or start from a look

Seedance 2.5: start and end frames, references and reference audio.

Rendering & your clips

While a clip renders you'll see **Queued** → **Generating** → **Finishing** and the time so far. Clips usually take 1–5 minutes.

- 1 **Not charged yet** — credits come off only when the clip lands. A failed or stopped clip costs nothing.
- 2 **Safe to close the tab** — it keeps rendering; come back and Studio picks the run up again.
- 3 **Use in a video** — opens a new Video Agent chat with the clip attached: “Make a video using this clip.”
- 4 **Reuse settings** and **Download** (mp4) on every row. Silent clips carry a *No audio* tag.



Seedance 2.0 Mini

Slow push-in on a ceramic coffee mug on an oak table by a window, steam rising, soft morning lightA ceramic coffee mug on an oak table by a window, soft morning light, ...

[Show all](#)

480p 4s 16:9 No audio

⚡ 40 credits

22 Sep 2026 · 01:37

Use in a video Reuse settings ⬇

A clip in History: model, shot, inputs, resolution, length, shape and price.

PART 4

Voiceover

Seven voice models, hundreds of voices — and your own. Seven lessons: script, price, voices, models, cloning, designing a voice and using what you make.

Your first voiceover

- 1 Open the **Voiceover** tab.
- 2 **Write the script** — type or paste what the voice should say. The price follows the character count as you type.
- 3 **Pick a voice** — press ▶ to hear any voice before you choose.
- 4 Press **Generate**. Most voiceovers arrive in seconds, saved to Files as an **mp3**.

CREDITS

The default, **ElevenLabs Multilingual v2**, costs 40 credits per 1,000 characters — a 135-character line is **6 credits**.

History

How it works

Make a voiceover in three steps

Every model, direct — and you see the credits before you press anything.

Meet the jacket that remembers every road you've taken

54 / 5,000 characters

priced as you type

R Roger
Laid-Back, Casual

S Sarah
Mature, Reassuring, Confid...

C Charlie
Deep, Confident, Energetic

▶

▶

▶

STEP 1

Write the script

Type or paste it. The price follows the character count as you write.

STEP 2


Pick a voice

Hear any voice before you choose. Voices you clone or design are in the list too.

STEP 3

Get your voiceover

A few seconds. It's saved to Files as an mp3 — download it, or let a clip on the Video tab follow it.



Don't know where to start?

Open Training — short lessons on prompts, looks and models.

>

Write & price the script

Under the script you'll always see **characters used / the model's limit** and the price in credits.

- 1 Priced by length** — credits per 1,000 characters, rounded up, with a minimum of 1 credit. Spaces at the start and end are free.
- 2 Limits** — 5,000 characters on ElevenLabs and Cartesia Sonic; 2,000 on Inworld, Kokoro and Chatterbox HD.
- 3 Over the limit?** Studio never cuts your script — it tells you how far over you are and waits for you to shorten it.
- 4** There are no speed or emotion sliders — pick the model and voice that fit the mood. **Eleven v3** is the expressive one.

TIP

Write for the ear: short sentences, numbers as words (“twenty-four seven”), and a comma wherever you want a breath.


The screenshot shows the Synthesys Studio interface. At the top is a large text input area for the script. Below it, a status bar displays '135 / 5,000 characters' and '6 credits'. Underneath, there's a model selection dropdown showing 'ElevenLabs Multilingual ...' with '6 credits · Text to Speech' below it. To the right of the dropdown is a large green 'Generate →' button.

A 135-character script on Multilingual v2: 6 credits, priced before you press Generate.

Pick a voice

The **voice** is who speaks; the **model** is the engine that speaks it. Click the voice card to open the list.

- 1 **Hear before you choose** — every ▶ preview is free.
- 2 **ElevenLabs** — a hand-picked library of 20 voices, like Roger, Sarah, Charlie and Alice.
- 3 **Inworld** — 113 voices in 15 languages, with a search box and language chips (English · 73, Spanish · 4...).
- 4 **Your voices** — clones and designed voices appear under **Your voices**, marked *YOUR CLONE* or *DESIGNED*. The bin beside one deletes it.
- 5 **Clone a voice** and **Design a voice** sit at the top of the list (lessons 4.5 and 4.6).



Left: the ElevenLabs voice library. Right: Inworld’s voices with language filters.

Voiceover models at a glance

MODEL	MAKER	WHAT IT OFFERS	CREDITS / 1,000 CHARS	MAX PER RUN	VOICES
Inworld TTS-1.5 Max	Inworld	Fast & cheap · 15 languages	2	2,000 chars	113
Kokoro	Hexgrad	Fast & cheap · American English	4	2,000 chars	20
Chatterbox HD	Resemble AI	Balanced · takes 2–4 minutes	8	2,000 chars	9
Cartesia Sonic	Cartesia	Balanced · 44 languages · your Studio voices	10	5,000 chars	Large library + yours
ElevenLabs Flash v2.5 ★	ElevenLabs	Balanced · fastest · 32 languages · cloned voices	20	5,000 chars	20 + your clones
ElevenLabs Multilingual v2 ★	ElevenLabs	Highest quality · 29 languages · the default	40	5,000 chars	20 + your clones
ElevenLabs Eleven v3 ★	ElevenLabs	Highest quality · expressive · 70+ languages	40	5,000 chars	20 + your clones

★ = Featured. The model list prices every row for the script you’ve written, so you can compare before you switch. Voices you clone or design in Studio speak on **Cartesia Sonic**; voices cloned in a Video Agent chat speak on the **ElevenLabs** models.

Clone your voice

Record or upload a short clip and Studio turns it into a voice you can use on **Cartesia Sonic**. Cloning is **free**.

- 1 Voice list → **Clone a voice**. Choose **Upload audio** (MP3, WAV, MP4 or WebM) or **Record audio**.
- 2 Give it **10–20 seconds** of clear, lively speech from one person. Only the first 20 seconds are used.
- 3 Name the voice, pick its language (15 to choose from), and leave **Remove background sound** on.
- 4 Tick the consent box — it must be your own voice, or you must have the speaker's explicit permission — then **Generate voice**.
- 5 **Up to 3 voices of your own** — clones and designed voices together. To free a slot, press the **bin** beside a voice under Your voices, then **Delete**.

TIP

A voice memo on your phone works great — hold the phone flat, about 6 inches from your mouth. When it's ready, Studio switches to Cartesia Sonic and selects your new voice.

Create voice clone

Speak with high energy, and record for at least 10 sec with a professional mic or smartphone.

Record audio

Upload audio

Drag and drop an audio or video file, or click to upload

MP3, WAV, MP4 or WebM

Pro tip: a voice memo on your phone works great.
Hold the phone flat about 6 inches from your mouth.

Create voice clone: record or upload, then name it and confirm consent.

Design a voice

No recording? Describe the voice you want and Studio makes three samples to choose from.

- 1 Voice list → **Design a voice**. Describe who is speaking: age, accent, mood — “a gravelly older man, slow and warm, like a late-night radio host.”
- 2 Stuck? Try a preset — Epic trailer narrator, Commercial (Positive), Corporate Narrator... — or the shuffle button for a surprise.
- 3 Press **Generate samples**: three takes of the same short English passage.
- 4 Play them, pick one, name it and **Save voice**. It appears under Your voices, marked *DESIGNED*.

CREDITS

11 credits per press of Generate samples (and per New samples). Saving is free. Samples are kept for 1 hour.

Design a voice

DESCRIBE THE VOICE109 / 1,000

A warm, unhurried woman in her 40s with a soft Irish lilt, like a late-night radio host telling you a secret.

✕

Epic trailer narrator

Commercial (Positive)

Inner Monologue

Conversational Narrator

Corporate Narrator

Loud Personality

3 samples · charged only when they arrive

Generate samples · 11 credits

Design a voice: describe it, or start from a preset.

PART 5

Tips & tricks

What separates a good result from a great one — and how to get there without wasting credits.

Prompts that work

- 1 **Image: subject → setting → light → lens → style.** “A frosted amber bottle on wet slate, soft window light from the left, 85mm, clean product photography.”
- 2 **Video: one shot, one move.** Say what moves and how the camera moves: “Slow push-in on the mug, steam rising, morning light sweeping across the table.”
- 3 **Voiceover: write for the ear.** Short sentences, numbers spelled out, commas where you want a breath.
- 4 **Describe people, don’t name them.** “A man in his 40s, salt-and-pepper beard” works; real names, brands and famous characters are the most common reason a model says no.
- 5 **Words in a picture?** Use **Ideogram v4** — it’s the model built for text inside the image.
- 6 **Let a Look write the first draft**, then edit the parts that matter to you.

Spend credits like a pro

- 1 Draft cheap, finish premium.** Find the idea with FLUX.1 [schnell] (1–3 credits) or a 480p MiniMax clip (25), then re-run the winner on a top model with **Reuse settings**.
- 2 Short first.** Prove a video idea at 4–5 seconds before paying for 15 or 30.
- 3 One still, many clips.** A start frame you like costs a few credits once — every clip you animate from it starts right.
- 4 ×4 only when you're exploring.** Four images cost four times one. Once you know what you want, go back to ×1.
- 5 Watch the shape.** On a few image models the aspect ratio changes the price — the chip always shows it.
- 6 Cheap voices for drafts.** Test a script on Inworld (2 credits per 1,000 characters) before the final take on ElevenLabs.


TIP

Nothing is charged until a result arrives, and failures are free — so it's always safe to press Retry.

From Studio to a finished video

Studio makes the pieces; the Video Agent puts them together.

- 1 **Image** → **clip** — make the key still on the Image tab, then pick it as the **start frame** on the Video tab.
- 2 **Voice** → **clip** — make a voiceover, then give it to a Seedance clip as **reference audio**.
- 3 **Anything** → **Video Agent** — **Use in a video** on any image or clip opens a chat with it attached; the agent builds the rest.
- 4 **Everything is in Files** — every Studio result is saved there, ready to attach anywhere in the app.



A clip made from a Studio still, playing in History.

PART 6

FAQ & troubleshooting

The questions everyone asks in week one.

Frequently asked

- 1 How long does it take?** Images: usually under a minute. Voiceovers: seconds (Chatterbox HD takes 2–4 minutes). Clips: usually 1–5 minutes.
- 2 Am I charged if it fails?** No. Credits are held when you press Generate and only come off when the result arrives.
- 3 Where do my results go?** History in Studio, and **Files** — automatically. Images and clips are also in the library's **Image generations** tab.
- 4 Can I close the tab?** Yes. A clip keeps rendering and is waiting in History when you come back.
- 5 Why can't I use my cloned voice here?** Voices made in Studio speak on **Cartesia Sonic**; voices cloned in a Video Agent chat speak on the **ElevenLabs** models. Switch the model and it's there.
- 6 How many voices can I keep?** Up to 3 of your own — clones and designed voices together. To make room, open the voice list and press the bin beside one under **Your voices**; it asks once more before it deletes.

When something looks off

- 1 **“The model wouldn’t make this one.”** Its safety filter said no. Describe people instead of naming them, drop brand names, use your own reference images — or try another model.
- 2 **Generate is greyed out.** Hover it for the reason — “Write a prompt first”, “Wait for the images to finish uploading”, “Shorten the prompt to fit”...
- 3 **My references disappeared.** The new model takes fewer — they’re held, not deleted, and come back when you switch to a model that takes them.
- 4 **My setting changed on its own.** The new model doesn’t offer it; Studio picked the nearest one and says what it changed under the bottom bar.
- 5 **“We lost the connection.”** The result may still arrive — check History before generating again.
- 6 **Something’s actually broken.** Every failure card carries a **Reference #** — quote it when you contact support.

TIP

Every failure card has **Retry** and **Try another model** — and neither costs anything until a result arrives.

SYNTHESYS STUDIO

Now go make something.

Every model, direct. The price before you press. Pay only for what arrives. — The Synthesys team ·
app.synthesys.io