



CREATOR STUDIO

UGC Studio v2.0

USER MANUAL V2

App: <https://aetherwavestudio.com/ugc-studio> (sign in required)

Reached from: the direct link above

Version: 2.0, September 2026

Every control named in this manual exists in the app today.

ICONIC! UGC Team the four founding members, each with a locked face, personality and voice



Katie Lee



Amy UGC



Guitar Rick



Cheryl Dear

AetherWave Studio · generated 2026-09-24

Written from the live code, not a spec. Every price below was read from the live pricing endpoint on 2026-09-24. Covers everything shipped through 2026-09-24: the v2.0 redesign (teams home, cast view, identity pack screen, 3-panel composer), multi-face forging, and draft saving.

■ The idea in one paragraph

UGC Studio builds a **team** of persistent AI creators for a brand, then has them make content. Each character locks three things: a **face** that stays the same in every post, a **personality** that writes every caption and script in their voice, and a **voice** for anything they say out loud. From a one-line brief, a character makes five kinds of post: a talking ad or product photo, a lifestyle photo, a candid behind-the-scenes video, a selfie vlog, or an opinion post. Every result saves to your gallery automatically, and one click hands it to Social-Cast for publishing.

■ What is new in v2.0

- **A full studio.** UGC Studio now sits inside the AetherWave platform (header, side menu, music player) and every screen is a three-panel studio.
 - **Teams home and cast view.** Your teams with their faces, a quick way back to the last character you used, and a profile panel for each character.
 - **The identity pack has its own screen**, with a live build log, a tick for each view, and a single **Approve and lock** moment.
 - **The Composer** puts who, what and the result side by side, shows the exact cost before you click, and keeps a strip of recent posts.
 - **Forge 1 to 4 faces at a time**, compare every face side by side, and download any of them.
 - **Your work is saved as you go.** An unfinished character survives a sign-out or a closed tab.
-

■ Part 1 - Teams and the cast

Teams home

Open `/ugc-studio` and you land on the teams home.

- **Left:** your teams, each with its cast faces and how many characters it has. **+ New team** creates one (it asks for a team name and an optional brand tag).
- **Middle:** a card per team with every member's face, **Open team** and **Compose**. Chips show how many identity packs are locked and how many characters still need one.

- **Click any face** to load that character into the right-hand column.
- The **Add member** tile at the end of the faces goes straight to forging a new character for that team.
- **Right: Pick up where you left off** shows the last character you made a post with, portrait-sized, with **Open composer** and **View profile**. After you click a face it reads **Ready when you are** and shows that character instead. Below it, **Recent across teams** shows your latest posts from every team.

The cast view

Open team shows one team's cast.

- **Left:** the team name, **Rename**, **Delete team**, and the cast list. A cyan dot means the character's identity pack is locked; a gold dot means it still needs one.
- **Middle:** a card per character: their face (the pack headshot once the pack is locked), the three pack views, their voice and vibes, and **Create post** (or a gold **Build identity pack** if there is no pack yet). The last card, **Forge a character**, starts a new one.
- **Right:** the selected character's profile: personality, how they speak, a voice sample you can play, their identity pack, and **Create post**, **Identity pack**, **Edit** and **Delete**.

⚠ **Deleting a team deletes every character in it.** Posts already made stay in your gallery.

■ Part 2 - Forge a character (4 steps)

A progress bar across the top shows the four steps.

Step 1: Who are they?

Fill in the identity fields: name, optional @handle, "They are a..." (woman, man, person), age, ethnicity, height and build, hair, face structure, eyes, lips, skin, and a distinguishing mark. Most fields offer suggestion chips; tap one to fill the field, or type your own. Together these become the **locked identity**, which is never rewritten after the character is created, and a live preview of it shows at the bottom.

- **Ethnicity:** specific beats broad. Typing "Korean" steers the face more reliably than "East Asian". Hover the ? for the category guide.
- **Distinguishing mark** is the single most useful field. One small unique detail, such as "subtle beauty mark under left eye", holds a face across generations better than more description.
- **Spellcheck** is on in every text field (name, handle and age are left alone).

Step 2: Lock the face

One neutral, front-facing image becomes the permanent reference. Neutral holds best: no smile, no sunglasses. Two tabs:

Forge it (recommended) generates studio portraits from the identity you wrote.

1. Pick **Faces per forge: 1, 2, 3 or 4**. The button shows the count and the cost, for example **Forge 4 faces - 36 cr**. Each face is 9 credits. The app remembers your choice.
2. Click **Forge**. The new faces appear at the top of the board as **Forging...** tiles, the page scrolls to them, and the button counts time (**Forging 4 faces... 0:32**). A face usually takes 20 to 90 seconds.
3. **Every face you forge stays on the board**, numbered and side by side, so you can compare a second or third round with the first. Click a face to pick it (pink ring and a check mark). The first face back is picked for you until you choose another.
4. Every face has a **Download** button, whether you use it or not.
5. Click **Use this face** to lock the one you picked.

Upload a photo uses an image you already have: drag a PNG or JPG onto the drop zone, or click it to choose a file (12MB max). You must tick the rights box ("I confirm I own the rights to this image and the likeness of any real person in it...") before **Use this photo** turns on.

Step 3: Give them a personality

This voice writes every caption, script and musing they post.

- **Vibe:** up to three of Funny, Deadpan, Warm, Hype, Chill, Nerdy, Bold, Wholesome, Chaotic, Luxe.
- **Backstory, Quirks + opinions, Speech style:** free text. The more specific, the more distinctive the captions.

Step 4: Pick their voice

Choose one of ten voices, each with a playable sample. It is used for talking videos and vlogs.

Voice	Style
Katie	Female, American English with a Chinese accent. Sweet, youthful, clear
Amy	Female, British. Pleasant, social-media native
Jessica	Female, American. Playful, bright, warm
Laura	Female, American. Enthusiastic, quirky
Sarah	Female, American. Reassuring, confident
Liam	Male, American. Energetic, social-media creator
Roger	Male, American. Laid-back, resonant
Eric	Male, American. Smooth, trustworthy
Brian	Male, American. Deep, comforting
Bill	Male, American. Wise, built for ad reads

Click **Create character**. You land on the new character's **identity pack** screen, because building the pack is the next step (Part 3).

Your work is saved as you go

Everything in the wizard, including every face you forged, is saved in your browser as you type. If you close the tab, reload, or get signed out, reopening **Forge a character** for that team brings it all back with a **Restored your unfinished character** bar. **Start fresh** clears it. If you are signed out mid-form, a bar says so, and its **Sign in again** link brings you straight back to the restored form. The saved copy is cleared once the character is created.

■ Part 3 - The identity pack

✓ **This is the most important step after creating a character.** The identity pack is three neutral views built from the locked face: **Front**, **Three-quarter** and **Full body**. Once locked, the pack *is* the face every post follows: photos and Backstage clips use all three views, and talking videos and vlogs animate the front headshot. Two or three matching views hold a face far better than one image, so a locked pack **replaces** the original image rather than riding alongside it.

Open it from **Identity pack** on a character's profile, **Build identity pack** on their card, or **Open pack** in the Composer.

- **Left:** the **source face** every view is built from, with **Replace face**.
- **Middle:** the three views and the pack controls.

- **Right:** the build log: a running clock, each view's state and time, and what locking means.

Building and locking

1. Click **Build identity pack**. The Front view builds from the source face; Three-quarter and Full body build from the Front view, so all three agree on one version of the person. Each tile shows its own clock while it builds, and the tiles still to come read **Queued**.
2. When all three are in, check each one and tick **Looks like [name]** under every view that reads as your character. If one is off, click that view's **Regenerate** to redo just that view.
3. **Approve and lock** turns on once all three are ticked. Click it. The views show **Locked**, the character's picture everywhere in the app switches to the Front headshot, and the cast list dot turns cyan.

Timing: about **6 minutes per view**, around **18 minutes** for a full pack. A long wait on one view is normal.

⚠ Keep the tab open while a pack builds. The views are built one after another from the page, so closing it stops the build. Views that finished are kept, and each one also lands in your gallery.

Cost: 9 credits per view, **27 credits for a full pack**. Regenerating one view is 9.

If it stops partway, the finished views are kept and the button reads **Resume building (N views left)**. Resume builds only the missing views, so you are never charged twice for a view you already have. Only one pack can build at a time.

Replacing the face

Tick **I have the rights to this image and the likeness in it**, then **Replace face** (up to 15MB). If the character already has pack views, the app asks you to confirm first.

⚠ Replacing the face clears the pack. A pack built from the old face would keep showing the old person, so you build a fresh one for the new face. This is the only way to change a character's face; the locked identity text stays read-only.

■ Part 4 - The Composer

Click **Create post** anywhere, or **Open composer** on the teams home. The Composer opens on the post type you used last with that character (Social Post the first time).

- **Left (who):** the team's cast, so you can switch characters without leaving (a gold dot means no pack yet); the character's face, vibes, voice and pack; and the five post types, each with its price.
- **Middle (what):** the brief for the chosen post type, then, pinned at the bottom, what it runs on, what it costs, how many credits you will have left, and **Generate**.

- **Right (the result):** a ready screen before you start, progress while it works, then the post, with the character's recent posts underneath.

The five post types

Post type	What you get	Price
UGC Ad	They pitch your product: a talking video in their voice, or a product-in-hand photo	Photo 9 credits. Talking video about 140 (see below)
Social Post	A photo of them living their life, with a caption	9 credits
Backstage Clip	A candid motion video, with spoken lines if you write them	Per second, from 24 credits (Part 5)
Vlog Entry	A selfie video of them talking to camera about their day	About 140 credits
Personal Musing	An opinion post in their voice, with a selfie	9 credits

Writing the caption and script is always free; you pay for the media.

Talking videos (UGC Ad talking, and Vlog Entry) cost 9 credits per second of video plus the voice (about 1 credit per 10 characters of script), so a typical 15-second read is roughly 140 credits. The price follows the script's length, so it is shown as an estimate. If the main lip-sync engine is busy, a backup engine takes over by itself; its rate can differ.

What each asks for

- **UGC Ad:** choose **Talking video** or **Product-in-hand photo**, then what they are selling, a call to action, and an optional product photo (**From my gallery** or **Upload**).
- **Social Post:** a scene idea. Describe clothes, light, place and action; hair and eye notes are ignored on purpose, because the face is locked.
- **Vlog Entry:** what the entry is about. Leave it blank and they riff on their day.
- **Personal Musing:** a topic. Leave it blank and they surprise you.
- **Backstage Clip:** see Part 5.

Photo posts (Social Post, Personal Musing and the product photo) also offer:

- **References:** up to 6 extra images (a location, an outfit, a product) **From my gallery** or **Upload**. They are used alongside the character's pack, never instead of it.
- **Format:** **Feed 3:4**, **Story 9:16** or **Square 1:1**.

Generating

The cost line above **Generate** shows the price and what you will have left ("1,275 left after, 9 cr"). If a fixed-price post costs more than you have, Generate is replaced by **Top up N credits to generate**, which opens the credits page in a new tab. A talking video's price is an estimate, so a short balance warns but does not block it.

While it works, the Generate button becomes a progress bar with a clock, and the right-hand panel shows a large timer, the current step, and the step list (caption, render, saving). **The caption appears as soon as it is written**, before the media is done. Once the final render has been handed off, the panel says **Safe to leave this page now** - the post finishes on our side and lands in your gallery. Before that point, keep the page open.

If something fails, the panel says what happened in plain words and offers **Try again** and **Edit brief**. Your brief is kept.

The result

- **The caption is editable**, with a character count against the 2,200 limit. Copy and Send with Social-Cast use your edited text.
- **Saved to gallery** is automatic; the button turns green.
- **Download** saves the file to your device.
- **Copy** copies the caption.
- **Send with Social-Cast** opens Social-Cast **in a new tab** with the caption and media loaded. The Composer stays open behind it.
- **Make another** returns to the ready screen with your brief kept.
- **Recent posts**: the character's latest posts sit under the result. Click one to see it again with its caption.

■ Part 5 - Backstage Clips (video)

Backstage is the post type that renders true motion video. It has its own controls.

Dialogue

Put lines in **"quotes"** in **What happens** and the character **says them on camera**. The engine creates the mouth movement and the voice together, so there is no lip-sync pass. With no quotes you get a wordless candid clip with ambient sound.

Example: `She looks up and deadpans, "We are doing this again? Really?"`

The take stretches to fit your lines. The app allows about 2.3 spoken words per second plus one second. If your lines need more time than the duration you picked, it extends the take to the next length that fits, and the price follows the longer take. The cost line shows this before you click Generate. If the lines need more than the engine's longest take, a warning says the end may rush.

Voice description

Voice description says how they sound, not what they say: pitch, pace, accent, energy. Using **the same description every time** keeps a character sounding like the same person across clips. It starts filled with the stock description for the character's voice, or whatever you last used for them. A **Sample** button appears when the text exactly matches a stock description; if you edit it, the button hides, because the sample would no longer be what renders.

Engines

Five engines, cheapest to most polished. Each card shows its live price per second at the resolution you choose. **Draft cheap, master expensive:** find the take on Grok or Mini, then re-run it on 2.0 or 2.5.

Engine	480p	720p	With a reference video	Lengths
Grok (budget draft)	4 cr/s	7 cr/s	not supported	6-15, 20, 25, 30s
Mini (clean motion)	6 cr/s	12 cr/s	8 / 16 cr/s	4-15s
Fast (standard, default)	17 cr/s	35 cr/s	22 / 48 cr/s	4-15s
2.0 (richer detail)	27 cr/s	58 cr/s	37 / 79 cr/s	4-15s
2.5 (premium)	40 cr/s	89 cr/s	48 / 107 cr/s	4-10, 12, 15, 20, 25, 30s

The defaults (Fast, 720p, 8 seconds) cost 280 credits. The same 8 seconds on Mini at 480p is 48. The cheapest possible clip is 24 credits. **Resolution** is 480p or 720p; **Aspect** is 9:16 (default), 16:9, 1:1 or adaptive. Typical render time is 1 to 3 minutes, or 2 to 4 on 2.0 and 2.5.

Director options

Open **Director options** for more control.

Start from is one switch with two choices, because the engine shoots in one mode or the other, never both:

- **References:** the character's face plus your images, composed into a new shot. Add **reference images** (drop them in, click to choose, or **From my gallery**), **reference videos** for motion or style (up to 3, 15 seconds total, 50MB each; using one switches to the reference-video rate), and **reference audio** (up to 3, 15 seconds total). Up to 9 images in total count toward the engine's limit, including the identity: **up to 8 extra** with a single face, **up to 6 extra** with a locked pack (the pack uses 3).
- **First frame:** one image that locks the exact opening shot. References are not used in this mode. On Grok, a first frame runs Grok 1.0. If you pick First frame without adding an image, Generate asks you to add one or switch back.

Switching back to **References** removes the first frame, so the two can never both be set. **Ambient sound** (on by default) controls audio on wordless clips; clips with spoken lines always carry audio. Grok takes image references only, and has no adaptive aspect (it uses 9:16 instead).

■ Part 6 - Managing characters

Edit (on a character's profile) changes the name, handle, vibe, backstory, quirks, speech style and voice. The locked identity text is read-only; the face changes only through **Replace face** on the identity pack screen, which clears the pack.

Delete removes a character. Their posts stay in your gallery.

Rename and **Delete team** are on the cast view.

■ Costs at a glance

Action	Credits
Forge a face	9 per face (1 to 4 per click)
Identity pack	27 (9 per view); regenerate one view 9
Caption or script	Free
Social Post, Personal Musing, product photo	9
Talking video or Vlog Entry	About 140 for a 15-second read (9 per second + voice)
Backstage Clip	Engine rate x seconds, from 24

■ Known gaps

- **No menu entry yet.** UGC Studio opens from its direct link, `/ugc-studio`. The grid menu's "UGC Ad" item opens Brand Studio's UGC mode, which is a different tool.
- **Recent posts and saved drafts live in your browser.** On another device or browser they start empty. Everything you made is always in your gallery.
- **The Backstage voice description is remembered per browser.** On another device it falls back to the stock description for the character's voice.
- **An identity pack builds from the open page.** Keep the tab open until all three views are in.
- **Resolution tops out at 720p** in UGC Studio, even though 2.0 renders 1080p elsewhere on the platform.
- **The backup lip-sync engine** can bill at a different rate from the main engine's 9 credits per second.

■ What changed from v1

- The whole app moved into the platform shell and three-panel studios: teams home, cast view, identity pack screen and Composer.
- Forging became 1 to 4 faces per click, with every face kept side by side and downloadable, and an upload drop zone.
- The identity pack became its own screen with a build log, per-view ticks and Approve and lock.
- The Composer gained a live cost line and top-up, a progress view with the caption shown early, plain-language errors with Try again, recent posts, photo formats, and gallery references.
- Unfinished characters are saved as you go and survive a sign-out.
- Spellcheck throughout the forms.
- Fixed along the way: identity packs for men and everyone else (the prompts once said "woman" for everyone); an honest reference-image limit with a pack; the pack headshot as the character's picture and talking face; hidden screens that stayed visible; and a pack that could have been saved onto the wrong character if you opened Edit on someone else mid-build.