

# THE FULL SYSTEM

## The AI Character Creator Kit

The complete system for building and running a consistent AI character brand — the same one used to build Cinna Bunny.

**\$27 · ONE-TIME PURCHASE**

### WHAT'S INSIDE

1. Why character-driven content works
2. Worked example — Cinna Bunny, fully filled in
3. The consistency brief framework
4. The prompt library — 16 lessons
5. Tools & stack reference
6. The production workflow — 16 lessons
7. Hook & caption formulas (7 formulas)
8. 30-day content calendar template
9. Metrics & what to track
10. Common mistakes & how to fix them
11. Growth & posting strategy
12. Turning this into revenue
13. FAQ

# 1. Why character-driven content works

Faceless AI content is everywhere now, which means it's also the easiest thing to scroll past. A character — a face, a personality, a running storyline — gives people a reason to come back. It builds parasocial attachment the same way a cartoon mascot or a recurring sitcom character does: viewers start following the character, not just the content, which is what turns a single viral clip into a channel people return to.

The trade-off is that a character has to look and feel like itself every single time, or that attachment breaks. This kit is built around solving that one problem, then using the same consistency to produce content efficiently at volume — from the first reference image through a full month of posting.

# 2. Worked example — Cinna Bunny, fully filled in

Before the framework itself, here's what it produces. Every field below is the exact, filled-in brief used to build and run Cinna Bunny — see Section 3 for what each field means and how to fill in your own.



*The consistency brief in practice — face, wardrobe, lighting and signature detail locked before a single frame is generated.*

|                  |                                                                                                                                       |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Core identity    | Cinna Bunny — anthropomorphic rabbit, stylized 3D, playful-glam young adult with a confident, slightly dramatic energy.               |
| Face             | Round, low-set eyes with a soft downward tilt; warm brown iris; soft blush cheeks; relaxed, faintly amused default expression.        |
| Body             | Dusty-rose fur, lighter chest and inner-ear patch, tall rounded ears, slim stylized proportions with slightly exaggerated ear length. |
| Signature detail | A single gold ear cuff on the left ear — present in every generation, regardless of outfit or scene.                                  |
| Home base outfit | Glam/nightlife default: sequins or metallic fabric, gold trim, tailored jacket silhouette.                                            |

|                           |                                                                                                                                     |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| <b>Wardrobe range</b>     | In-brand: glam, formal, seasonal glam variants. Explicitly not: sporty, casual streetwear, anything that reads "generic AI avatar." |
| <b>Colour palette</b>     | Dusty rose, gold, deep charcoal — recurs across fur, outfits, sets, and lighting grade.                                             |
| <b>Lighting</b>           | Warm gold rim light with soft ambient fill; key light slightly above eye level.                                                     |
| <b>Lens / framing</b>     | Medium shot, slight low angle, shallow depth of field to keep focus on the character.                                               |
| <b>Colour grade</b>       | Warm highlights, deep shadow, slight film grain for texture.                                                                        |
| <b>Personality / tone</b> | Confident, a little dramatic, sassy captions, never meek or apologetic.                                                             |
| <b>Never do</b>           | Sporty poses, apologetic captions, muted/desaturated colour grades, off-brand casual wardrobe.                                      |

That's the entire brief — twelve short fields, reused in every single prompt — and it's what makes the character read as "the same bunny" across 50+ generations and multiple content types.

### 3. The consistency brief framework

This is the expanded version of the brief. Every field you lock in here gets pasted, in shorthand, into every prompt you write — for images and for video. Treat it as your character's bible, not a one-time exercise.

#### Core identity

- Name, species/form, age and general vibe in one sentence.

#### Physical reference

- Face: eye shape, size, colour, expression defaults.
- Body: proportions, texture, base and secondary colour.
- One or two signature details that are unmistakable at a glance (an accessory, a marking, a colour combination) — these are what let viewers recognise the character even in a thumbnail.

## Style & wardrobe

- A default "home base" outfit used whenever a scene doesn't call for something specific.
- A defined wardrobe range — what's in-brand and what's explicitly not (this prevents drift into generic AI-character territory).
- A 2–3 colour palette that recurs across outfits, sets, and lighting.

## Camera & world rules

- Lighting setup (key, rim, ambient).
- Lens and framing defaults (shot distance, angle, depth of field).
- Colour grade (highlight/shadow tone, grain, contrast).

## Personality (for captions, voice, and story)

- Tone in 3–4 words.
- A short "never do" list — poses, expressions, or lines that would feel out of character.

Keep this brief in a single doc or note you can copy from every time. The goal is that any image or video generation you produce could be dropped into the existing feed without looking like a different character.

## 4. The prompt library — 16 lessons across 5 content types

Five content types cover almost everything a character-driven account needs. This section breaks each one into a lesson on when to use it, a lesson on the exact template structure, and a lesson on applying and varying it — 16 lessons in total, starting with how the library fits together.

### Lesson 1 — How the prompt library works

Every template below is built to receive your consistency brief from Section 3 as a drop-in block. The brief never changes; only the scene, action, and content-type-specific structure around it changes. Read the template, fill the bracketed fields, paste your brief where indicated, and generate. This is what keeps five very different content types looking like the same character.



*One locked brief, five content types — intro shots, POV/transformation, story-arc, lifestyle b-roll, seasonal.*

### Content type A — Character intro / reveal shots

#### Lesson 2 — When to use it

Used for profile content, pinned videos, and any clip introducing the character to a new viewer. Static, hero-style framing that shows the character clearly. This is usually the first piece of content a new viewer sees, so it carries the most weight for first impressions.

#### Lesson 3 — The template, field by field

##### Template

*[Character, full consistency brief] standing in [signature setting], [default outfit], [lighting setup], [lens/framing], confident expression, looking directly at camera, [colour grade].*

#### Lesson 4 — Applying it and varying it

Keep the pose simple and the framing static — this shot's job is clarity, not motion. Reuse the same signature setting across intro shots so the location itself becomes recognisable, and pin the strongest result to the top of the profile.

## Content type B — POV / transformation clips

### Lesson 5 — When to use it

The character reacting to something, transforming, or revealing a look — built for the hook-in-the-first-second style of short-form video. Best used when you want a scroll-stopping opener rather than a slow build.

### Lesson 6 — The template, field by field

#### Template

*[Character, consistency brief] in [before state], sudden [transformation/reveal action], ending in [after state / outfit], dynamic motion blur, [lighting], quick cut framing, [colour grade].*

### Lesson 7 — Applying it and varying it

The "before state" should contrast hard with the "after state" — the bigger the visual gap, the stronger the hook. Pair this template with the transformation cut hook formula from Section 7 for the strongest combination.

## Content type C — Story-arc scenes

### Lesson 8 — When to use it

Recurring characters, relationships, or an ongoing storyline (this is what "Cinna & Roxy"-style content is built on). These clips work because viewers come back to see what happens next — treat it like a running series, not one-off content.

### Lesson 9 — The template, field by field

#### Template

*[Character, consistency brief] and [second character, their own consistency brief] in [scene/setting], [interaction/conflict/moment], [emotional tone], [lighting], [lens/framing], [colour grade].*

### Lesson 10 — Applying it and varying it

Give every recurring second character their own consistency brief, filled in with the same rigour as the lead. Plan story beats a few clips ahead so each scene answers a question the last one raised — that's what keeps comments asking "what happens next."

## Content type D — Glam / lifestyle b-roll

### Lesson 11 — When to use it

Lower-stakes filler content — the character simply existing in their world. This is what keeps a feed feeling alive between bigger story or hook clips, and it's the cheapest content type to produce at volume.

### Lesson 12 — The template, field by field

#### Template

*[Character, consistency brief] in [lifestyle setting — nightlife/home/travel/etc.], [default or seasonal outfit], candid pose, [lighting], [lens/framing], [colour grade].*



Batch these — generate five or six in one sitting across different settings so you always have filler ready between higher-effort posts. Vary the setting more than the outfit to keep it feeling fresh without breaking wardrobe consistency.

## Content type E — Seasonal / trend-jacking content

### Lesson 14 — When to use it

Timely content tied to a holiday, trending audio, or a format everyone's using that week — the character's consistency brief keeps it feeling on-brand even when the concept is borrowed.

### Lesson 15 — The template, field by field

#### Template

*[Character, consistency brief] participating in [seasonal moment / trend concept], [seasonal outfit variant within wardrobe range], [lighting], [lens/framing], [colour grade], energy matching the trend's tone.*

### Lesson 16 — Applying it and varying it

Only borrow trends that fit inside the defined wardrobe range from Section 3 — a trend outfit that reads off-brand does more damage to recognition than the trend is worth chasing.

## 5. Tools & stack reference

This kit is tool-agnostic by design — the framework works with whatever you already use or plan to use. Here's how the pipeline typically breaks down by category, so you know what to look for at each stage.

### Reference & image generation

- Any AI image generator capable of consistent character/reference-image workflows (image-to-image or character-reference features are the key capability to look for).

### Video generation

- Any AI video generator that accepts a reference image as a starting frame — this is what carries your locked-in look from Section 2 into motion.

### Voice & audio

- AI voice generation for character dialogue, or royalty-free/trending audio matched to the personality tone from your brief.

### Editing

- Any short-form editor with caption and pacing tools — mobile editors work fine for this format; the workflow doesn't require anything specialised.

Swap tools freely as better ones come out. What matters is that every tool in the chain receives the same consistency brief and reference image, not which specific brand you use.

## 6. The production workflow — 16 lessons

The pipeline is tool-agnostic — swap in whichever AI image/video generator, voice tool, and editor you already use. The order and the checkpoints matter more than the specific software. This section walks the full pipeline as 16 lessons: the seven production steps, four quality checkpoints, and five lessons on running it efficiently week over week.



*The production pipeline, start to finish — from locked brief to a posted, on-brand clip.*

### Lesson 1 — Workflow overview

The pipeline runs in a fixed order for a reason: each step locks in something the next step depends on. Skipping or reordering steps is the single most common cause of a character drifting off-brand mid-production.

### Lesson 2 — Step 1: lock the consistency brief

Lock the consistency brief before generating anything. Every field from Section 3 should be written down, not just remembered — a brief that only exists in your head drifts within a few generations without you noticing.

### Lesson 3 — Step 2: generate a reference batch

Generate a small batch of reference images first (5–10) using the brief. This batch is where you catch a face, colour, or proportion that isn't reading the way you intended, before it propagates into video.

### Lesson 4 — Step 2b: selecting your anchor image

Pick the 1–2 references that most closely match the intended look — these become your visual anchor for everything after. Save them somewhere permanent; you'll return to this exact pair every time the character starts to drift.

### **Lesson 5 — Step 3: generate video from the anchor**

Generate video from the approved reference images, using the relevant prompt template from Section 4 for the content type. Always generate from the saved anchor image, never from a previous video frame — chaining video-to-video is the fastest way to drift.

### **Lesson 6 — Step 4: add voice or audio**

Add voice/audio if the clip needs it, matching the personality tone defined in the brief. A mismatched voice or audio energy undercuts a technically perfect visual just as fast as a wrong outfit does.

### **Lesson 7 — Step 5: caption with a hook formula**

Caption using one of the hook formulas in Section 7. Match the formula to the content type — a story-arc scene pairs naturally with the relationship tease, a transformation clip with the transformation cut.

### **Lesson 8 — Step 6: edit for pacing**

Edit for pacing — the first 1 second carries the most weight; cut anything before the hook lands. If the hook isn't visible or audible in the first frame, move it earlier rather than trimming around it.

### **Lesson 9 — Step 7: post and track**

Post on a consistent cadence (Section 11) and track which content type and hook style is performing, then weight future production toward what's working — see Section 9 for exactly what to track.

### **Lesson 10 — Quality checkpoint: face & signature detail**

Does the face, proportions, and signature detail match the reference image? Compare side by side, not from memory — small drifts are easy to miss one generation at a time but add up fast across a week of output.

### **Lesson 11 — Quality checkpoint: wardrobe range**

Does the outfit fall inside the defined wardrobe range? An outfit that's technically well generated but outside the range still reads as off-brand to a returning viewer.

### **Lesson 12 — Quality checkpoint: lighting & colour grade**

Does the lighting and colour grade match the brief, not just "look nice"? A generation can look great in isolation and still break continuity if it doesn't match the grade used in surrounding content.

### **Lesson 13 — Quality checkpoint: caption tone**

Does the caption tone match the personality field? Tone drift in captions is the most common checkpoint failure because it's easy to write "in the moment" rather than checking it against the brief.

### **Lesson 14 — Batching for efficiency**

Run steps 2 and 3 in batches rather than one clip at a time — generate a week's worth of reference images and video passes in a single sitting, then move through voice, captioning, and editing as a second pass. Batching cuts total production time significantly compared to finishing one clip fully before starting the next.

### **Lesson 15 — Troubleshooting drift mid-pipeline**

If a generation starts drifting, don't push forward and fix it in editing — go back to Step 2's approved anchor image and regenerate from there. Drift compounds the further it travels down the pipeline, so the earliest fix is always the cheapest one.

## **Lesson 16 — The weekly production rhythm**

Tie this pipeline into the 30-day calendar from Section 8: batch references and video early in the week, layer voice/captions/edits mid-week, and leave the final day for quality checkpoints and scheduling the coming week's posts.

## 7. Hook & caption formulas

The same principle that makes an ad hook work applies to a character clip: you have about one second before a viewer decides to keep watching. Seven formulas that consistently work for story/character content:



*Seven hook formulas that consistently earn the first second — and the caption tone that pairs with each.*

- **The mid-drama drop-in** — open mid-conflict or mid-reaction with no setup ("POV: you asked for the longer version and she showed up") — forces the viewer to watch to understand what's happening.
- **The bold statement** — a single confident, slightly provocative line as the opening caption or spoken line — sets tone immediately and filters in the right audience.
- **The relationship tease** — reference an ongoing storyline or second character by name, rewarding regular viewers and creating curiosity for new ones.
- **The transformation cut** — text on screen naming what's about to happen ("after the fall"), then delivering the payoff in the next 1–2 seconds.
- **The direct address** — character speaks or reacts straight to camera, breaking the fourth wall — this consistently outperforms passive/observational framing for watch-through.
- **The relatable admission** — the character says or captions something mildly self-deprecating or overly honest — humanises the character and invites comments.
- **The countdown / stakes** — frame the clip around a ticking clock or consequence ("3 days until the reveal") — gives viewers a reason to come back for the payoff.

## 8. 30-day content calendar template

A sample first week is filled in below to show the rhythm — repeat and adjust the pattern for the rest of the month based on what's performing (see Section 9).

| Day | Content type         | Hook formula            | Notes                                             |
|-----|----------------------|-------------------------|---------------------------------------------------|
| Mon | Intro / hero shot    | The bold statement      | Anchor post for the week — pin if new to profile. |
| Tue | Lifestyle b-roll     | The relatable admission | Low-effort filler; keeps feed active.             |
| Wed | Story-arc scene      | The relationship tease  | Advance the ongoing storyline.                    |
| Thu | POV / transformation | The transformation cut  | Highest-effort clip of the week.                  |
| Fri | Lifestyle b-roll     | The direct address      | Casual, conversational tone.                      |
| Sat | Story-arc scene      | The countdown / stakes  | Cliffhanger going into the weekend.               |
| Sun | Seasonal / trend     | The mid-drama drop-in   | Tie to whatever's trending that week.             |

Repeat this weekly rhythm for the full 30 days, swapping in new scenes and settings while keeping the mix of content types consistent. Two story-arc slots per week is usually enough to build continuity without overwhelming production capacity.

## 9. Metrics & what to track

Track these weekly, not daily — daily numbers are too noisy to act on. Look for trends across at least 5–7 posts before changing strategy.

- **Watch-through rate by content type** — this tells you which of the five formats (Section 4) is earning full views, not just impressions.
- **Follow rate per post** — views that convert to follows matter more than raw view count.
- **Comment themes** — are people asking "what happens next," naming the character, or requesting more of a specific storyline? That's your growth signal.
- **Hook formula performance** — tag each post with which formula from Section 7 it used, and compare watch-through rates across formulas over time.
- **Drop-off point** — if most viewers leave at the same timestamp, that's usually a pacing or hook problem, not a content problem.

## 10. Common mistakes & how to fix them

|                                                                                  |                                                                                                                                                              |
|----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ✗ <b>Changing small details "just this once" because a generation looks good</b> | Reject the generation and regenerate with the brief instead. A one-off drift is exactly what starts undermining recognition.                                 |
| ✗ <b>Treating every post as one-off content</b>                                  | Build in story-arc slots (Section 8) even at low volume — two per week is enough to create the "what happens next" pull.                                     |
| ✗ <b>Posting inconsistently, then judging performance</b>                        | Give any new posting cadence at least 2–3 weeks before drawing conclusions — algorithms and audiences both need a consistent signal.                         |
| ✗ <b>Vague brief fields ("cute," "pink")</b>                                     | Rewrite with specifics: exact shade, exact proportion note, exact signature detail. Vague fields are the #1 cause of visible drift.                          |
| ✗ <b>Skipping the reference-image step</b>                                       | Always generate and approve a small reference batch (Section 6, Lesson 3) before jumping to video — video generation without an anchor image drifts fastest. |

## 11. Growth & posting strategy



*Cadence and content mix, working together — the pattern behind Cinna Bunny's growth.*

- Post on a consistent cadence rather than sporadically — regularity is what turns casual viewers into followers who expect new content.
- Mix content types deliberately: intro/hero shots build recognition, POV/transformation and story-arc clips drive views and follows, lifestyle b-roll keeps the feed feeling alive between bigger moments.
- Build continuity — recurring characters, callbacks, and running storylines give people a reason to come back, the same mechanic that makes serialized content sticky.
- Read comments for "what happens next" energy and use it — the audience will tell you which storylines and characters to keep investing in.
- Revisit your metrics (Section 9) every 1–2 weeks and shift the content mix toward whatever's earning the strongest watch-through and follow rate.

## 12. Turning this into revenue



*From tips to done-for-you — the monetization ladder for a character-driven account.*

A few paths, roughly in order of effort:

- **Tips** (Ko-fi, TikTok gifts) — low effort, low ceiling, works best once a small dedicated fanbase exists.
- **A shop of your own** — merch, digital wallpapers, or prompt packs like this one, sold directly to the audience the character has already built.
- **Brand deals** — once the account has reach, brands will pay for a sponsored appearance or crossover.
- **Done-for-you** — if building and running this yourself isn't the goal, KBV Publishing builds custom AI characters for brands end-to-end: consistency brief, content pipeline, and ongoing production. Cinna Bunny is a live example of the same system covered in this kit, applied at scale.

**Want it built for you?** Book a free strategy call and we'll map out the right done-for-you package for your brand live — bring your product and goals.

**Book your free DFY strategy call →**

**Prefer to build it yourself, with support?** Join creators running this exact system inside the free AI Unchained community — swap prompts, get feedback, and see what's working right now.

**Join the free community →**

## 13. FAQ

### **Do I need to use the exact tools you use for Cinna Bunny?**

No — the framework in this kit is tool-agnostic. It works with any AI image generator, video generator, and editor that supports a reference-image or image-to-image workflow. See Section 5 for what to look for in each category.

### **How long before a character-driven account starts growing?**

Give any new posting cadence at least 2–3 weeks of consistent output before judging results (Section 10). Most drift and performance issues in the first few weeks trace back to an under-specified brief, not the platform or the tools.

### **Can I run more than one character at once?**

Yes — each character needs its own consistency brief (Section 3) filled in separately. Story-arc scenes (Section 4C) work well for introducing a second character alongside an established one.

### **What if my character starts looking different after several generations?**

Go back to your reference images from Section 6, Lesson 4, and regenerate using the original approved reference plus your written brief — not a recent generation, which may have already drifted.

### **I want this built for me instead of doing it myself — is that an option?**

Yes. KBV Publishing offers a done-for-you service covering the consistency brief, content pipeline, and ongoing production, the same system used to build and run Cinna Bunny. See Section 12 for details, or **book a free strategy call** to map out the right package.

### **Is there somewhere to get feedback or ask questions as I go?**

Yes — the free AI Unchained community is where creators using this system trade prompts, troubleshoot drift, and share what's actually working right now. [Join for free →](#)