

THE PLAYBOOK

Build your own AI design agent.

Stand up a small AI worker that turns your copy into **on-brand visuals** — the same slides every time, no drifting fonts or off-brand colors. Here's the pipeline I actually use to design my posts.

by Corina · The CV Method · @the_cv_method



WHAT YOU'RE BUILDING

Not “ask AI to make me a pretty slide” — that drifts every single time, new fonts, new colors, a different vibe. Instead: **lock your brand into a fixed template, and let the agent only fill in the words.** The design is decided once; the agent pours approved copy into slots it can't move or restyle. A hundred posts come out looking like one brand — because they're the same template, every time. It rests on one rule.

RULE ZERO

The brand is the source of truth. The agent fills the words — it doesn't redesign.

The agent never invents styling. It maps your approved copy into a locked template — same layout, palette, and fonts every run. That single constraint is what makes it consistent instead of a slot machine.

Four pieces. Here’s what I use and the general equivalent — swap in whatever you already have.

PIECE	WHAT IT DOES	I USE	SWAP FOR
The brand system	Your colors, fonts, and spacing in one place the template reads.	A tokens file (CSS variables)	A Canva Brand Kit · Figma variables
The template	A fixed layout with named text slots (cover · content · stat · CTA).	An HTML template kit	Canva templates · Figma frames
The renderer	Turns a filled template into an image at exact size.	A headless-browser screenshot	Canva/Figma export
The agent	Maps your approved copy → the slot text, one entry per slide.	Claude + a small script	Any tool-use LLM

88 STEP 2 · LOCK THE BRAND INTO TOKENS

02

DECIDE THE LOOK ONCE

Put the brand in one place

Define your palette, fonts, and spacing as **tokens** in a single file the template reads — say a navy base, a green accent, one bold pop color, and your three fonts. Every template pulls from these, so nothing is decided ad-hoc.

Change a token, and every future post updates with it. This “decide it once” step is what actually kills the drift — the agent literally has no other colors to reach for.

Keep **one gradient moment** per slide, and never put body text on a live gradient — it reads as brand polish and stays legible.

03

THE AGENT FILLS WORDS, NOT LAYOUT

A handful of slide types, each with slots

Design a few slide types by hand — **cover, content, stat/quote, CTA** — and give each one **named text slots** (headline, body, footer). That's it. The agent's only job is to put the right words in the right slots.

Because it can't move elements, choose fonts, or recolor anything, it **can't go off-brand** — the worst it can do is write weak copy, which you catch on review. Design the fence once; pour words into it forever.

Four to six slide types cover almost everything. Add a new type deliberately, not once per post — that's how the system stays tight.

04

STRUCTURED, NOT FREEFORM

Copy in, a slide spec out

Hand the agent your **approved** caption. It returns a simple **spec**: which template each slide uses, and the text for every slot — one entry per slide. Plain and small enough to read at a glance before anything renders.

Two honest rules ride here: any client proof uses a **category label**, never a real name; and a stat slide carries **real numbers only** — the agent never invents a figure to fill the space.

Keep the spec human-readable (a little JSON or a table). Eyeballing the spec is faster than re-rendering — you fix wording before you spend a render.

05

PIXEL-EXACT, WITH A HUMAN GATE

Render, look, then post

Render each filled slide to an image at the **real dimensions** (1080×1350 for a carousel), so what you see is exactly what the feed shows. Then **you review the images** — the human gate — and post the ones that land.

The agent drafts the visuals; you approve them. Same as every good AI worker: it proposes, you ship. Nothing posts itself.

Rendering is deterministic — the same spec makes the same slides — so what you approve is precisely what goes out. No surprise re-interpretation.

STEAL THESE PROMPTS

COPY → SLIDE SPEC – HAND THE AGENT YOUR APPROVED CAPTION

Here's my approved caption: **[paste]**.

Turn it into a slide spec for my template. For each slide, output: template type (cover / content /

stat / cta) and the text for each named slot (headline, body, footer).

Rules: use ONLY my brand template – don't invent styling; 5–8 slides; category labels for any client

proof (no real names); real numbers only (never invent a figure to fill space).

Return it as a simple list I can eyeball before I render.

KEEP IT HONEST & ON-BRAND

The **brand system** is the source of truth

On-palette only · one gradient moment

Real numbers only — no invented stats

Category labels, never a client's name

Approved logos, clear-space respected

You approve every image before it posts

THE STARTER CHECKLIST

- ☐ Put the brand in **tokens** — colors, fonts, spacing.
- ☐ One **gradient moment** per slide; no body text on a gradient.
- ☐ **Category labels** for proof; **no invented numbers**.
- ☐ **Review every image** — the human gate.
- ☐ Build **4-6 fixed slide templates** with named slots.
- ☐ Agent maps approved copy → a **slide spec** (template + slots).
- ☐ Render at **exact size** (1080×1350 carousel).
- ☐ Post — the agent drafts visuals, **you approve**.

Want a branded design agent — or the whole content system — built for you?

I set up the tokens, the templates, and the render pipeline so your posts look like one brand every time — then hand you the keys.

hello.thecvmethod.com