// CASE STUDY · AI CONTENT PIPELINE

Dick Tracy Origins.

Building character-consistent AI video from a public-domain source. Four frames, one Character Pack, three modular environment plates, zero IP exposure.

// FIGURE 01 · EPISODE 1 CHARACTER PACK · FOUR IDENTITY-LOCKED ANCHORS, GENERATED AS 3D FIGURINE REFERENCES BEFORE ANY SCENE WORK

PROJECT · DICK TRACY ORIGINS, EPISODE 1

SOURCE · 1931 TRUEHEART ROBBERY STRIPS

OUTPUT · 40 SECONDS OF FINISHED VIDEO

COST · $21.50 IN STACK-TIME, OR $0.54/SECOND

↓ executive summary

The audience can spot drift in *one frame.*

Generative AI can produce a single beautiful image. The harder problem is producing a hundred of them that share the same face, the same coat, the same room, and the same light. This proof of concept solved that problem for a four-frame opening sequence drawn from the 1931 origin of Dick Tracy, a property whose earliest strips lapsed into the public domain.

We built the sequence end to end using the Nano Banana Professional Framework (NBPF), a five-stage pipeline designed around Google's Gemini 3 Pro image model. Characters were generated once, in isolation, and then composited into modular environment plates. Lighting, wardrobe, and set continuity held across all four frames. The work was governed by an explicit IP framework that renamed protected marks while keeping every visual element legally clean.

The full LarryDoJo production stack (Claude Max, Google AI Ultra, and Adobe Premiere) runs $472.98 a month. This POC consumed one working day of that stack and produced forty seconds of finished, character-consistent video. That works out to $21.50 in stack-time, or fifty-four cents per finished second. The result is a repeatable production method.

HEADLINE OUTCOMES

- → Forty seconds of finished video produced from four narrative starting frames, all anchored to a single Character Pack and three modular environment plates.
- → Character consistency held across frames sharing a location (Frames 1 and 3) without manual retouching.
- → IP exposure reduced to zero: the protagonist was renamed Detective Hawknose while preserving every public-domain visual element.
- → Cost per finished second of video came in at fifty-four cents.
- → Workflow documented stage by stage so a second team can repeat it without supervision.

40s	$21.50	$5.38	$0.54
Of finished video, four hero frames.	One day's slice of the production stack.	Per finished hero frame, fully composited.	Per finished second of video. Less than a stamp.
// DELIVERABLE	// STACK-TIME	// PER FRAME	// PER SECOND

↓ contents

What's *inside.*

↓ 01 · the problem

The frame is fresh *every time.*

Generative video, in early 2026, has a continuity problem. Every frame is, by default, a fresh draw. A character’s nose shifts. A coat changes shade. A room rearranges itself. For a single hero image this does not matter. For a story, it matters completely.

The team set out to test whether a structured, stage-gated pipeline could deliver narrative-grade consistency from a commodity image model. The brief had three constraints:

01 Use only public-domain source material.

02 Produce a recognizable, period-correct opening sequence.

03 Document every step so the workflow is repeatable.

We chose the opening of the Dick Tracy strip, October 14 to 15, 1931. The strip is a clean test case: a small cast, two locations, a clear emotional arc from warm domestic scene to violent intrusion, and a copyright that lapsed decades ago for failure to renew.

⚠ WHY THESE STRIPS?

The 1931 Trueheart Robbery introduces everything we need in four panels.

Protagonist, love interest, family, antagonists. A warm interior, a cold exterior, and a doorway that bridges them. For testing character and environment continuity, the source could not be better chosen.

↓ 02 · source & ip strategy

Public domain follows a *slow release.*

Copyright on the 1931 strips lapsed in the late 1950s when the Chicago Tribune failed to renew. The underlying 1931 imagery (the canary-yellow overcoat, the black fedora, the square-jawed detective) is therefore free to use. The trademark on the name "Dick Tracy" is not.

Trademarks do not expire so long as they remain in commerce. Tribune Content Agency continues to license the name. Using it in a published title or as a brand mark would invite a takedown regardless of the underlying art's public domain status.

Our solution: rename the character, keep the look

We renamed the protagonist Detective Hawknose. The name is descriptive, it nods to Chester Gould's original silhouette, and it carries no enforceable mark. Every visual element of the character (the coat, the hat, the jaw, the slicked hair) comes directly from the 1931 strips and is therefore in the clear.

// SAFE AS-IS	
Tess Trueheart	PUBLIC DOMAIN
Emil Trueheart	PUBLIC DOMAIN
Mrs. Trueheart	PUBLIC DOMAIN
Crutch & partner	PUBLIC DOMAIN
Yellow overcoat & fedora	PUBLIC DOMAIN

// NAMES & ELEMENTS LOCKED	
"Dick Tracy" name	TM ACTIVE
Flattop (1943)	UNTIL 2039
2-Way Wrist Radio (1946)	UNTIL 2042
Moon Maid (1964)	UNTIL 2060
1990 film likeness	RIGHTS HELD

Element-by-element clearance

We logged every visual element against the year it first appeared. A famous Tracy villain like Flattop did not appear until 1943 and stays locked until 2039. The 2-Way Wrist Radio (1946) stays locked until 2042. We used neither. The 1931 cast of villains (Big Boy, Stooge Viller, Crutch) we used freely.

Clearance log

ELEMENT	FIRST APPEARANCE	STATUS	USED IN POC
Yellow overcoat & fedora	1931	Public domain (failed renewal)	Yes
Tess Trueheart character	1931	Public domain	Yes
Emil & Mrs. Trueheart	1931	Public domain	Yes
Crutch (villain)	1931	Public domain	Yes
Name "Dick Tracy"	1931	Active trademark	No (renamed)
Flattop	1943	Locked until 2039	No
2-Way Wrist Radio	1946	Locked until 2042	No
Moon Maid	1964	Locked until 2060	No

↓ 03 · the framework

Five stages. *Gated.*

The Nano Banana Professional Framework breaks production into five gated stages. Each stage produces a defined output, and no stage begins until the previous one validates. The pipeline is explicit on purpose: it removes guesswork and makes the work auditable.

> ⚠ THE CARDINAL RULE
>
> **Never generate a character directly inside a scene from a text prompt.**
>
> Independent research measured a consistency score drop from 7.99 to 0.55 when this rule is broken. Always generate the character anchor first, in isolation, then composite into the environment.

The five stages at a glance

STAGE I	**DNA Definition** Write the Character Bible and generate the 3-view Character Pack.	IDENTITY-LOCKED PACK
STAGE II	**Plate Construction** Build modular environments as reusable kits.	3 PLATE KITS
STAGE III	**Anchor Visualization** Generate characters in isolation, composite into plates.	IDENTITY-LOCKED FRAME
STAGE IV	**Temporal Chaining** Use the last frame as reference for the next.	CHAINED SEQUENCE
STAGE V	**Semantic Refinement** Edit with natural language; never re-roll.	FINAL FRAME

↓ 04 · stage i

Character *DNA.*

The first stage answers a single question: what does this character look like, in enough detail that the model can reproduce them on demand? The answer takes two forms, a written Character Bible and a visual Character Pack.

The Character Bible

Each principal received a written specification covering anatomy, materials, and fixed identifiers. The fixed identifiers matter most. They are the verbatim phrases repeated in every prompt, "canary-yellow wool overcoat, double-breasted with six black buttons", that the model treats as anchors.

⚠ SAMPLE BIBLE ENTRY

Detective Hawknose

Male, early 30s. Square jaw with a strong chin cleft. Dark hair, slicked back with pomade, visible comb lines. Sharp, narrow eyes, deep-set. Prominent aquiline nose. Wearing a canary-yellow wool overcoat, double-breasted with six black buttons. Black fedora with a yellow grosgrain ribbon band. White cotton dress shirt, black silk tie with a Windsor knot.

The Character Pack

From each bible we generated a three-view Character Pack (front, three-quarter side, and back) in a single Nano Banana Pro session to lock internal consistency. We rendered the characters first as 3D collectible figurines. The plasticine geometry of a figurine is easier for the model to reproduce than a fully rendered human face, so figurines became the foundation that later, more photorealistic versions were built on top of.

// FIGURE 02 · TESS TRUEHEART AND DETECTIVE HAWKNOSE · WE TESTED TWO HAIR COLORS FOR TESS BEFORE SETTLING ON THE WARMER RED TO CONTRAST AGAINST HAWKNOSE'S DARK HAIR AND THE YELLOW COAT 08

Validation: the Exploded View test

Once the Pack was rendered, we ran the Exploded View Validation. This is a stress test: we asked the model to deconstruct the character into individual clothing layers with three head close-ups showing different emotions. If the model can perform this disassembly while keeping the face and proportions stable, the anchor is production-ready. Hawknose passed on the second attempt; the first run drifted on the coat’s button count and we tightened "six black buttons" into the bible as a fixed identifier.

// FIGURE 03 · DETECTIVE HAWKNOSE, HERO RENDER · THE YELLOW COAT IS THE VISUAL ANCHOR OF THE ENTIRE PRODUCTION

↓ 05 · stage ii

Plates are *kits.*

A plate is not a finished scene. It is a kit. The discipline of Stage II is to build the world as reusable parts (walls, floors, lighting, props) that can be recombined for different camera angles without losing continuity.

Episode 1 needed three plates:

PLATE	LOCATION	MOOD / LIGHTING	USED IN
A	Trueheart living room interior	Warm tungsten, modest, domestic	Frames 1, 3, 4
B	Rooftop exterior at night	Cold moonlight, high-contrast noir	Frame 2
C	Trueheart doorway transition	Harsh cold burst (composited into A)	Frame 4

PLATE A · THE WARM BASELINE

The living room had to carry three different frames, including the violent reversal of Frame 4. We built it warm and lit it from a single overhead pendant at 2700K, with a standing lamp providing fill and a window adding faint blue-gray ambient. The lighting was generated warm, on purpose. The cold-light version for Frame 4 came later through Stage V semantic editing, not a fresh build.

PLATE B · THE COLD COUNTERWEIGHT

The rooftop served one purpose: to feel like the opposite of Plate A. Cold moonlight from upper-left, deep impenetrable blacks, a single warm rectangle of light on the parapet underside representing the Trueheart window seen from above. That warm rectangle is the most important detail in the plate. It tells the audience, without dialogue, that the figures on the roof are watching the family below.

PLATE C · THE SEAM

Plate C is the smallest and most demanding. It exists to be composited into Plate A at the moment the door bursts open. Material continuity matters absolutely here: the floral wallpaper, the dark oak floor, and the doorframe molding had to match Plate A pixel for pixel. We generated Plate C in the same Nano Banana Pro session as Plate A, with Plate A loaded as a reference image, to force material continuity.

↓ 06 · stage iii

Anchor *visualization.*

Stage III is where the work is won or lost. Character anchors and environment plates are combined here, and a single shortcut taken at this stage cascades into every later frame. We followed the Cardinal Rule strictly: every character was generated independently in a neutral gray studio, validated against the Character Pack, and only then composited into the plate.

JSON blueprinting

Before any image was generated, each frame received a structured JSON blueprint. The blueprint defines camera focal length, framing, lighting setup, character positions, poses, expressions, and the target plate. It is the architectural plan the model follows during compositing. We generated the blueprints with Gemini 3 Reasoning before handing them to Nano Banana Pro.

The four frames

FRAME 1 · EMIL COUNTING MONEY

Emil sits at the dining table counting a thousand dollars in cash. Mrs. Trueheart stands behind him, hand on his shoulder. The family is debt-free for the first time. Lighting: warm pendant pool, standing lamp fill, window ambient. The mood is hope and pride. The calm before the storm.

FRAME 2 · THE SURVEILLANCE

Two figures crouch on the rooftop opposite, peering down at the lighted Trueheart window. They have spotted the money. We rendered them as silhouettes against the cold sky, with only a faint cold rim catching the brim of a fedora and the curve of a shoulder. Their identities are not yet revealed to the viewer. They are threats, not people.

FRAME 3 · THE ENGAGEMENT ANNOUNCEMENT

Hawknose and Tess stand together, arms linked. Emil sits in the green velvet armchair, looking up at them. Hawknose has just asked for Tess’s hand and Emil has given his blessing. This is the emotional peak of the domestic scene, and the first frame in which the protagonist’s yellow coat appears. The coat had to be the warmest, most saturated object in the frame. The pendant light made the yellow wool glow.

// FIGURE 05 · EMIL ISOLATED RENDER · ARMCHAIR POSE FOR FRAME 3

// FIGURE 06 · FRAME 4 SETUP DETAIL · COLD AMBIENT PREVIEWS THE REVERSAL

IP-Adapter tuning

For each composite we set the IP-Adapter scale between 0.4 and 0.6. The scale balances the strength of the character anchor against the influence of the plate’s lighting and texture. For hero introductions like Frame 3 we used 0.55. For silhouetted villains in Frames 2 and 4 we dropped to 0.4, since silhouette work needs body shape rather than facial fidelity. Starting Control Step stayed at 0.4 throughout, injecting character identity early in the generation while letting the plate finish the lighting.

FRAME	IP-ADAPTER SCALE	RATIONALE
Frame 1	0.50	Balanced warmth and identity
Frame 2	0.40	Silhouettes need shape, not face detail
Frame 3	0.55 / 0.50 / 0.50	Hawknose highest: hero intro frame
Frame 4	0.40	Backlit silhouettes, body shape only

↓ 07 · stage iv

Temporal *chaining.*

Once the four starting frames were locked, we chained them. The Temporal Bridge rule is simple: the last frame of any generated sequence becomes the reference frame for the next. This preserves the dust in the light cone, the position of the money on the table, the angle of the candelabra. Without chaining, the world resets between cuts.

The chaining map for Episode 1

01 Frame 1 locks the warm baseline.

02 Frame 2 chains style only (the Character Pack), but uses a new plate (Plate B). The mood reverses but the world remains coherent.

03 Frame 3 chains directly from Frame 1: same plate, same lighting, Emil moves from the table to the armchair. The continuity must be exact because the audience reads any drift as a mistake.

04 Frame 4 chains from Frame 3, then composites Plate C into the background and reduces pendant intensity in the foreground. The reversal lands because the warm baseline was held perfectly until the moment it broke.

> ⚠ WHY THIS MATTERS
>
> **The dramatic power of Frame 4 depends on how steady Frames 1 and 3 were.**
>
> If the warm baseline drifts, the cold reversal feels like a different movie. Continuity is not aesthetic polish. It is the load-bearing wall of the narrative.

↓ 08 · stage v

Semantic *refinement.*

No first composite is final. Stage V applies natural-language edits to the locked frames (changing an expression, shifting a light, repositioning a prop) without re-rolling the whole image. The rule: if a frame is 80% correct, edit; do not regenerate.

Frame 4 took the most refinement. We needed to reduce pendant intensity to sell the mood shift, fine-tune the warm-to-cold gradient on the floor where Plate C’s light spilled in, and adjust silhouette rim light. Five passes brought it to final. Frames 1 and 3 needed two passes each, mostly to balance Emil’s expression against the warmth of the pendant.

What we did not edit

We resisted the temptation to edit faces. Once a face composited cleanly from the studio render into the plate, we left it alone. Editing a face under semantic masking risks identity drift; the model can subtly reshape features while it is "just adjusting the smile." The discipline is to fix the lighting, fix the pose, and trust the anchor.

EDIT TYPE	RISK LEVEL	STRATEGY
Lighting change	Low	Environment edits rarely affect identity
Expression change	Medium	May alter facial geometry slightly
Clothing swap	Medium-High	Can trigger material drift
Character addition	High	May destabilize existing anchors

↓ 09 · governance

Quality *gates.*

Every stage passed through a written checklist before the next began. The checklists are not decoration. They are the difference between a one-off success and a repeatable process.

GATE	QUESTION	IF IT FAILS
I.1	Does the Character Bible specify anatomy, materials, and fixed identifiers?	Tighten the bible; do not generate.
I.2	Did the Character Pack pass Exploded View Validation?	Tighten fixed identifiers; regenerate.
II.1	Are the plates rendered empty, with no characters?	Regenerate; this is the Cardinal Rule.
II.2	Do shared materials match across plates?	Regenerate Plate C with Plate A as reference.
III.1	Did each character generate independently before compositing?	Stop. Re-do; consistency will collapse.
III.2	Did the studio render match the Character Pack?	Re-render with tighter prompt.
IV.1	Does each frame chain from the previous one’s reference?	Re-chain. Do not skip ahead.
V.1	Did identity drift during semantic editing?	Roll back; re-edit with tighter prompt.

Documentation discipline

Each frame was archived with three artifacts: its JSON blueprint, the studio renders used to compose it, and the prompt history. Anyone returning to this work in six months can rebuild any frame from these three files alone. We treat the blueprints as source code: version-controlled, reviewed, and never edited in place.

↓ 10 · tools & cost

Forty seconds. *$21.50.*

A real production stack costs real money. The honest cost of this POC is not "API tokens." It is one day of three professional subscriptions: the reasoning layer, the image and video layer, and the editing layer. Together they produced forty seconds of finished, character-consistent video.

The monthly stack

Claude Max 20x · reasoning, blueprints, orchestration	**$200.00 / mo**
Google AI Ultra · Gemini 3 Pro, Nano Banana Pro, Veo	**$249.99 / mo**
Adobe Premiere Pro · final edit and assembly	**$22.99 / mo**
Total monthly stack	**$472.98 / mo**

The math

$472.98 / mo ÷ 22 working days	**$21.50 / day**
$21.50 / day ÷ 4 finished hero frames	**$5.38 / frame**
$21.50 / day ÷ 40 seconds of video	**$0.54 / second**

Fifty-four cents a second. A traditional animation house quotes 60 seconds of character animation at $5,000 to $25,000. The math problem the framework solves is not "how do we make AI cheaper than humans." It is "how do we make AI consistent enough to be worth using at all." Once that's solved, the price drops out of the equation.

↓ 11 · what we learned

Six things *that mattered.*

// 01

The Cardinal Rule is real.

We tested, briefly, what happens when you generate a character directly inside a scene. The drop in consistency is immediate and obvious. Faces shift, coats lose buttons, fabric textures change. Stage III exists for a reason.

// 02

Fixed identifiers are everything.

The phrases that survived from the Character Bible into every generated prompt did the most work. Generic descriptions ("yellow coat") drifted. Specific descriptions ("six black buttons") held.

// 03

Plates are kits, not paintings.

Treating an environment as a single finished image is a trap. Treating it as a kit of recombinable elements is what lets the same room hold three different frames without breaking.

// 04

Edit, do not re-roll.

Every time we re-rolled a frame at 80% rather than editing it, we lost something. A small detail, a lighting balance, a pose we liked. Stage V exists to preserve the wins.

// 05

IP discipline pays for itself.

Renaming the protagonist took an afternoon. The legal certainty it bought us is worth more than every other safeguard combined. Every project should start with a clearance pass, not end with one.

// 06

The framework is the product.

The frames themselves are pretty. The framework is what makes them repeatable. A second team given the bibles, the blueprints, and a Gemini 3 Pro account could reproduce this work in a day.

↓ 12 · next steps

What ships *next.*

This proof of concept established that character, set, and style can hold across a four-frame sequence. The next phase tests the framework at video length and at scale.

PHASE 2 · ANIMATE THE FOUR FRAMES

Each starting frame becomes a 5 to 8 second video clip using image-to-video with start-end frame guidance. The Temporal Bridge from Stage IV extends naturally into video: the last video frame of clip N becomes the first video frame of clip N+1. We expect the warm-to-cold reversal of Frame 4 to be the most demanding test of the bridge so far.

PHASE 3 · EXTEND TO A FULL EPISODE

Episode 1 of Dick Tracy Origins runs roughly two minutes of screen time. That requires twelve to sixteen starting frames, two new plates (the police precinct and a Chicago street exterior), and one new character anchor (Chief Brandon, who appears in the second half of the source strips). The framework absorbs this without modification. We add bibles, we add plates, we run the same five stages.

PHASE 4 · PACKAGE THE FRAMEWORK

The pipeline is documented. The next deliverable is a packaged framework: templates for the Character Bible, the JSON blueprint, and the gate checklists, plus a reference repository of the Episode 1 outputs. With those in hand, a second team should be able to start a new episode without supervision.

⚠ THE BIGGER QUESTION

What does production look like a year from now?

If one person can produce forty seconds of publication-grade episode opener for $21.50 in stack-time, fifty-four cents per finished second, what does production look like a year from now? The framework solves the consistency problem. The remaining problem is taste. And taste is exactly the part that does not scale by adding more compute.