

Every campaign has deliberate cross-asset differences and accidental ones.

Canva has no way to distinguish between them.

THE GAP BRAND KIT DOESN'T COVER

Brand Kit enforces visual consistency — same colors, fonts, logos across any design. That layer exists and works.

But a campaign isn't one design. It's five to fifteen assets — landing page, email, social posts, ads, blog — edited by different people across different days. Some differences are *correct*: Instagram says "Try It Free" because that converts better on mobile. LinkedIn says "Get Started" because that matches professional tone. Those are deliberate platform adaptations.

The problem: accidental differences look identical to deliberate ones. And they multiply with every person and every edit.

Brand Kit handles visual identity. Templates handle layout reuse. Neither gives a campaign a single source of truth for its own content — the price, the CTA, the launch date. That's the gap.

SEE IT HAPPEN — ONE CAMPAIGN

FitPulse "Summer Shape-Up" — 4 people, 5 assets, 2 weeks

7 elements shown. A real campaign has **20+ shared elements × 8–15 assets × 6 markets**.

ELEMENT	LANDING PAGE	EMAIL	INSTAGRAM	LINKEDIN	BLOG
CTA	Start Free Trial	Start Free Trial	Try It Free	Get Started	Start Free Trial
Price	\$9.99/month	\$9.99/mo	\$9.99	\$9.99/month	\$12.99/month
Date	June 1	June 1st	—	June 1	June 2
Headline	Get Fit This Summer	Summer Shape-Up Starts Now	Your Summer Glow-Up	Transform Your Routine	Get Fit This Summer
Offer Details	30-day free trial	Free trial included	—	30-day free trial	14-day free trial
Accent Color	Summer Orange	Summer Orange	Summer Orange	Brand Teal ←	Coral
Promo Code	SUMMER24	SUMMER24	—	SUMMER24	SUMMER23

- Error — wrong data
- Adaptation — deliberate
- Ambiguous — who knows?

Red cells — price updated day 3, blog never synced. "Brand Teal" = defaulted to Brand Kit because there's no place for campaign colors without polluting it. **Blue cells** — platform CTA conventions, not errors. **Amber cells are the core problem** — a reviewer can't tell if four headlines are intentional platform messaging or four people freelancing without reading the brief. That brief lives in a Google Doc nothing in Canva has ever read.

~2,400 cells

A mid-size team runs 12 campaigns/quarter × 10 assets × 20 shared elements. That's ~2,400 cross-references to audit — **per quarter**. Nobody maintains that spreadsheet. A customer screenshots both prices → fire drill.

77–81%

of company content ships with brand guideline violations even when guidelines exist. *"Files live in different places, reviewers weigh in too late or not at all, and high-risk work is submitted without the right checks."* That's visual consistency — content consistency (CTAs, dates, pricing) has no tooling at all.

Marq, 400+ orgs surveyed · Canva Design Operations Guide

See everything. Know what's intentional. Fix what's not.

THE CONSISTENCY TAXONOMY

Not all differences are errors. The taxonomy defines which ones matter.

● Must-Match

Factual content that must be identical — if it differs, it's a bug

Prices, dates, URLs, promo codes, disclaimers

● Platform-Adapted

Intentional variation the AI recognizes and does NOT flag

Instagram "Learn More" vs. landing page "Start Free Trial", shorter mobile copy

● Should-Match

Style worth reviewing — AI surfaces it, human decides

Tone, hero image, color emphasis

Without this taxonomy, the scanner flags everything and teams ignore it. With it, the scanner knows Instagram CTAs *should* differ and that pricing *should not*. This is our v1 hypothesis — calibrated during beta.

THE FULL FLOW

1 Set Up Campaign Kit

Campaign-scoped kit: brief, canonical values, and campaign-specific overrides (seasonal palette, co-brand colors) — so teams don't have to pollute Brand Kit or ignore its warnings. AI extracts values automatically.

2 Review Adaptations

AI suggests platform-specific variations (e.g., Instagram CTA → "Learn More"). Team approves or dismisses. Approved adaptations won't trigger false flags.

3 Scan

AI checks every text element against Campaign Kit + Brand Kit + approved adaptations, then compares across all assets simultaneously. Results streamed in real time (SSE).

4 Fix

Each issue shows severity, affected assets, and the recommended value. One click to propagate across every affected asset. Preview per asset. Full undo.

Campaign Kit → Canonical Values → Scanner Context. Without a kit, the scanner uses majority voting. With one, it has ground truth — an asset that disagrees with the kit is wrong even if every other asset also disagrees.

Where This Lives: A new **AI Projects** section in Canva Teams — each project is a campaign with its assets, Campaign Kit, and health score. Brand Kit = permanent, org-level visual identity. Campaign Kit = ephemeral, campaign-scoped content identity (brief, canonical values, campaign fonts/colors, approved adaptations). Campaign ends, kit goes away. Brand Kit never changes.

Live Prototype

eulerity-design-studio.vercel.app

Claude API + SSE streaming. 21 editable campaign assets. Live canvas editor with drift detection. Natural language fix ("make all CTAs match the brief"). Guided onboarding tour.

Measure what matters. Kill what doesn't move.



NORTH STAR METRIC

Cross-Asset Consistency Rate at Publish

Baseline: est. ~60–70% of campaigns ship with undetected issues (no tool exists to measure — baseline established in Phase 1 beta). Target: **95%+** consistency at publish (directional; threshold set after baseline). Kill condition: if <30% of campaigns show issues in Phase 1, problem isn't big enough for Phase 2.

WHAT WE MEASURE

SYNC ALL ADOPTION

>60%

of must-match issues resolved via one-click fix over manual editing.

TIME TO CONSISTENCY

<2 min

per campaign (up to 5 assets, must-match). Fast enough to become a pre-publish habit.

FALSE POSITIVE RATE

<20%

must-match flags. Measured via user dismissal rate. Over-flagging preferred over missed errors.

REPEAT USAGE

>50%

campaign-over-campaign. Teams come back because it solves a real problem.

WHAT WE HYPOTHESIZE (VALIDATED POST-LAUNCH)

- Post-publish inconsistency rate drops >80% from baseline established in Phase 1
- Cross-asset scanning becomes a top-cited reason for Teams tier adoption in post-upgrade surveys
- Campaign throughput increases as consistency checking shrinks QA cycles

⚠ **Guardrail:** Any confirmed incorrect propagation is a **PO** that blocks rollout. Sync All error rate stays below **1%**. An incorrect propagation is worse than no propagation.



24 Weeks. 4 Phases. Each gated on results.

THE CORE Weeks 1–10: Ship the product.	
1: Detect Weeks 1–4	<i>"It caught the wrong date on LinkedIn. I didn't notice."</i> Campaign workspace + text scanner + Consistency Cards. Detection only — proves accuracy before we build fixes. ♦ Gate: >80% precision
2: Fix Weeks 5–10	<i>"I synced the CTA across 5 assets in one click."</i> Sync All for text + visual consistency + format-dependent intelligence. ♦ Gate: <1% Sync All error rate
THE SCALE Weeks 11–24: Make it infrastructure.	
3: Scale Weeks 11–16	<i>"We haven't published with a wrong date in months."</i> Auto-scan on save + enterprise admin controls + team dashboard. ♦ Gate: >90% consistency at publish (active adopters)
4: Learn Weeks 17–24	<i>"I can't imagine launching without scanning first."</i> Pre-publish gates + cross-campaign learning + multi-language. Prevents drift, not just catches it. ♦ Gate: >40% fewer post-publish corrections

Each phase gated on results. No phase starts until the previous one proves its value. Teams involved: **Engineering** (2–3 engineers, canvas API integration), **Design** (1 designer, consistency cards + fix flow UX), **AI/ML** (taxonomy model calibration, prompt engineering), **Data** (baseline measurement, metrics instrumentation), **Enterprise Sales** (beta partner recruitment). Assumes existing canvas API access.

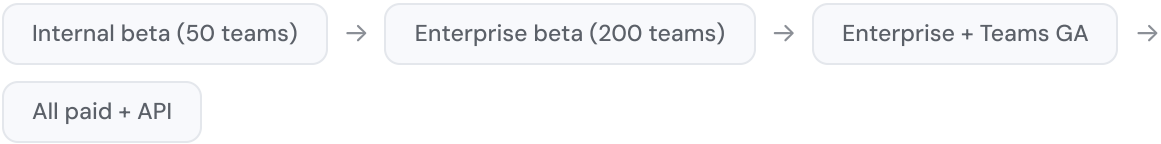
WHAT WE'RE NOT BUILDING

- × **Not a PM tool.** No timelines, briefs, or task assignment. That's Monday/Asana — would bloat scope to 32+ weeks.
- × **Not auto-approve.** No enterprise lets AI publish unsupervised. Scanner surfaces problems. Humans decide.
- × **Not a brand compliance checker.** Brand Assist handles palette compliance. We handle content consistency — different category entirely.

KEY RISKS

- × **Brief adoption.** Most teams don't have a single canonical brief. Majority-voting fallback works without one, but accuracy degrades.
- × **AI classification accuracy.** Misclassifying a platform adaptation as must-match erodes trust. Taxonomy calibration is the Phase 1 priority.
- × **Canva API dependency.** Write access to canvas elements required for Sync All. Not publicly available today — requires platform team partnership.
- × **LLM cost at scale.** Claude API call per scan × campaigns/quarter. Mitigation: cache canonical values, batch scans, model distillation in Phase 3+.

ROLLOUT





6 Directions Evaluated. 5 Cut. 1 Shipped.


The hardest part wasn't writing — it was deciding what to build. Every direction got a dedicated adversarial Opus agent that stress-tested its weakest assumptions. The reframe that created the Consistency Agent came from a question I asked: *"But why would it shift if they are using Brand Kit?"* — which exposed that Brand Kit handles palette consistency, not content consistency. That single challenge killed the original direction and reshaped everything.

THE DECISION FUNNEL

✗	Feedback Fusion AI clusters reviewer comments into insight cards. Strip away the AI and it's a comment filter. Organized human input, didn't generate new intelligence.
✗	Brand Memory AI learns implicit design patterns over time. Fatal cold start: needs months of history before it's useful. No team waits that long.
✗	Brand Review Agent Evolve Brand Assist into scoring + auto-approve. Canva already has Brand Assist. Auto-approve is anti-collaborative: brief says "teams working together," not "AI replaces the review step."
✗	Campaign Pilot Full end-to-end campaign orchestrator: brief intake, assignments, timelines, consistency. Ballooned to 32-week scope. The PO features were single-player. 70% overlapped with Monday/Asana.
✗	AI Design Partner Real-time creative coaching copilot. Fails the solo-user test — works identically for one person as for a team. Canva is already building this with their AI Assistant.
✓	Consistency Agent Cross-asset content consistency. The gap is real (no tool checks whether your Instagram CTA matches your landing page), it gets worse at scale, and it only works as a team feature.

RESEARCH VALIDATION — "DID CANVA ALREADY FIX THIS?"

Two parallel Sonnet agents cross-referenced every pain point against Canva's actual product changelog (2024–2026). If a problem was already solved, the direction was dead.




Sources: Reddit, G2, Capterra, Canva Community, Canva Help, LinkedIn, YouTube, HubSpot, Miro — 40+ URLs. **Data:** 77–81% brand violations (Marq, 400+ orgs), 61% need 3–5 review rounds (Adobe), 20–30% of marketing time wasted on asset tasks.

Agent usage: 8 Sonnet research agents · 6 Opus adversarial agents · 4 Opus spec agents · 4 Opus build agents · 3 Opus review agents. Full log: [ai-workflow-log.md](#)

How the Prototype Got Built in 3 Days

Four parallel workstreams — research, critique, code, deployment — each handled by a different AI tool.

THE PIPELINE



THE KILL PROMPT — ADVERSARIAL, RAN AGAINST EVERY DIRECTION

"Find the weakest assumption. What would a PM with 10 years of experience see through immediately? If this is the obvious answer, tell me now."

THE FACT-CHECK PROMPT — CAUGHT AN AI ERROR IN THE DECK

AI claimed "false negative rate <5%" as a metric. Devil's advocate found it indefensible — no ground truth dataset exists to measure false negatives. Removed. Replaced with taxonomy calibration via user dismissal rate.

DIVISION OF LABOR

AI handled **volume** — 40+ sources, 6 parallel evaluations, 14.9K lines of code, deployment automation. The kill decisions on Slide A, the taxonomy design, and the "Brand Kit handles visual identity, not content consistency" reframe were mine. The log has every prompt.

The full AI workflow log — every prompt, every pivot, every kill decision — is in **ai-workflow-log.md**.