WHITE PAPER · UPDATED SEPTEMBER 2026

AI WON'T SAVE YOUR BRAND. YOUR OPERATING MODEL WILL.

A field guide for brand and marketing leaders who are tired of being told to “embrace AI.” Includes AI-stack playbooks, a creative-operations risk briefing and a governance checklist.

Travis Hurlock

Creative Direction & Brand Strategy

travishurlock.com

EXECUTIVE SUMMARY

READ THIS PAGE FIRST

CMOs are already spending on AI. Most aren't seeing it back. This paper argues the gap isn't the tools — it's the operating model around them — and sets out what to do about it before your next AI budget decision.

70% of CMOs admit their organisation isn't mature enough to scale the AI they've already bought (Gartner, 2026)	95% of generative AI pilots inside large organisations show zero measurable P&L impact (MIT, 2025)	4X more likely a consumer distrusts a brand after spotting AI content than trusts it more (Klaviyo/ YouGov, 2025)

Chapters 1–3 make the diagnostic case: the readiness gap, why bolt-on AI adoption fails while redesigned workflows succeed, and the trust cost of visible AI content. **Chapter 4** shows what a properly directed AI stack looks like for three real creative-operations tasks. **Chapter 5** gives the four moves to make first, in order. **Chapter 6** is the one most reports skip — the risks specific to running AI through a creative and marketing operation, from shadow AI to IP exposure to brand impersonation, with what to actually do about each. **Chapter 7** shows how this maps onto real client work already running today.

If you read nothing else: read the pull-quotes. Each one is the whole argument of its section, compressed to a sentence.

INTRODUCTION

THE REPORT NOBODY NEEDED ANOTHER OF

You've already adopted AI. So has everyone in your category. That was never the hard part.

Seventy percent of CMOs will admit, if you ask them off the record, that their teams aren't ready to use the AI budget they've already signed off. Not because the tools don't work. Because nothing *around* the tools has changed — same approval chains, same briefs, same definition of what “good” looks like.

This isn't another report telling you to embrace AI. There are enough of those, and none of them are working. This is about the four things separating the brands getting real value from it, the stacks actually worth building, the risks specific to a creative operation that a general IT security policy won't catch, and the first move to make before your next AI decision.

01 THE READINESS GAP

CMOs now put 15.3% of their marketing budget toward AI. Only 30% say their organisation is actually mature enough to scale what they've bought.

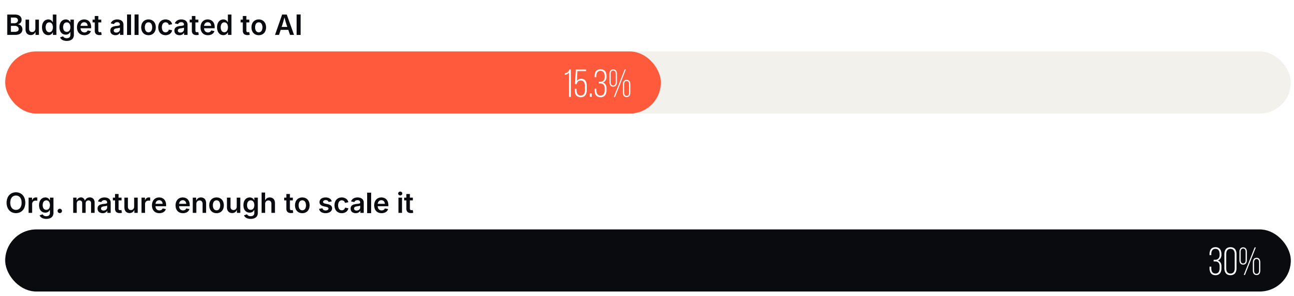

Source: Gartner, 2026 CMO Spend Survey.

That gap isn't a technology problem. It's an operating model problem wearing a technology costume. You can hand a team the sharpest tool in the market, and if the brief still takes three weeks to approve, the tool just makes you wait faster.

Budget was never the constraint. Structure is.

If the tool is faster than the decision-making around it, the tool isn't the bottleneck anymore. You are.

02 THE ADOPTION TRAP

Here's the number that should worry you more than any AI headline this year: 95% of generative AI pilots inside large organisations show zero measurable impact on profit or loss.

The ones that *do* work share one trait — they weren't bolted onto the existing process. They replaced a step of it.

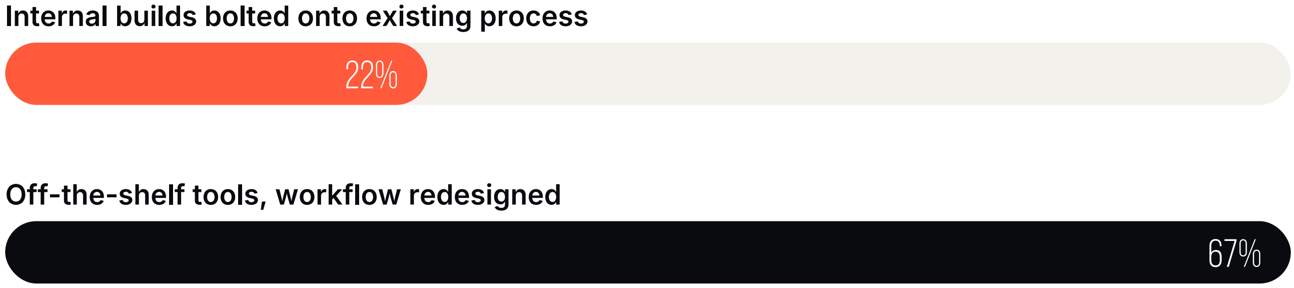

Source: MIT NANDA / Initiative on the Digital Economy, July 2025. Figures describe pilot success rates, not universal outcomes.

Picture the brand team that bought an AI copy tool to "speed up social." Six months on, drafts arrive faster — and then sit in the same four-person approval chain they always did, waiting the same five days. Output went up. Nothing that mattered changed. That's not an AI failure. That's a refusal to touch the process the AI was supposed to improve.

Bolting AI onto a broken workflow doesn't fix the workflow. It just gives it a faster engine and the same brakes.

03 THE TRUST TAX

Consumers are more than four times as likely to actively distrust a brand after spotting AI-generated content as they are to trust it more because of it. Over half say they're uncomfortable seeing visibly AI-made content on brand social channels at all.

Trust the brand more

7%

Distrust the brand

31%

Source: Klaviyo/Datalily consumer study, Dec 2025, with YouGov.

Meanwhile, inside the industry, the creative director's job is quietly being rewritten — from the person who makes the work to the person who decides which of a hundred AI-generated options is worth showing anyone. As one global creative lead put it recently: getting clicked isn't a measure of performance. That's a measure of a click.

Speed was never the scarce resource. Judgment is.

> **The brands that win this decade won't be the ones producing the most. They'll be the ones still capable of saying no.**

04 AI STACKS FOR CREATIVE OPERATIONS

Every AI stack article makes the same mistake: it lists tools and calls that a strategy. The stack is never the strategy. Who's directing it is.

Below is the operating model this paper is written from — six senior roles, each responsible for a different judgment call, with AI tooling doing the acceleration underneath each one, not replacing the decision itself.

01

The Strategist

Brand strategy, positioning and messaging. First read on every brief before a pixel gets touched.

02

The Creative Director

Visual direction and concept development. Holds the line on taste and consistency across every output.

03

The Copywriter

Campaign copy, tone of voice and scripts, built to a defined voice, not a generic one.

04

The Analyst

Data interpretation, reporting and insight, so decisions are made on evidence, not instinct alone.

05

The Social Lead

Content planning and platform strategy, matched to where the audience actually is.

06

The Systems Architect

Workflow design, tool stack and automation. The plumbing that lets the other five scale.

Three real briefs, and the stack actually built for each:

A Repositioning a brand that's stopped making sense

Leads: The Strategist

Claude | OpenAI | Figma | Adobe CC

Output: New positioning, messaging framework and identity system — the “Clarify the proposition” brief.

B Launching something that has to land on day one

Leads: The Creative Director + Social Lead

Midjourney | Runway | Framer | Mailchimp | Semrush

Output: Launch creative, a digital experience built to convert, and a sequenced campaign — the “Launch or grow something important” brief.

C Scaling content without losing the standard

Leads: The Systems Architect

Notion | Asana | Zapier | ElevenLabs | Metricool | Google Analytics

Output: A governed production workflow with a judgment gate before anything publishes — the “Fix how creative work gets delivered” brief.

Same discipline every time: a named senior lead accountable for the brief, AI tooling doing the acceleration, and a human check before anything reaches a customer. Swap the tools in five years. Keep that structure.

The stack is never the strategy. It's just how fast the strategy gets made real.

The gap this paper is actually about

Adoption has outrun oversight, and the sector's own numbers show it. Marketing teams have raced to put AI into production — almost none have put governance in at the same pace.

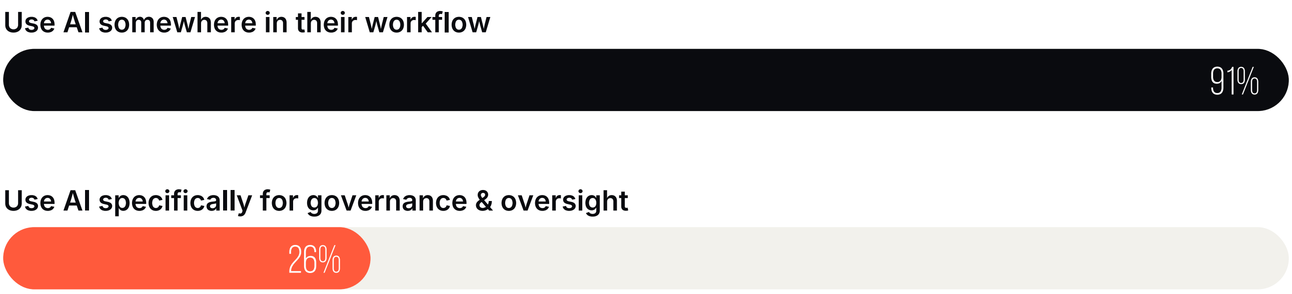

Source: Jasper, State of AI in Marketing 2026.

98% of organisations with an AI governance platform in place report a positive return on their AI investment.	34% of organisations without one do. Same tools, same budgets — the gap is entirely the operating model around them.

Source: SAS & Coleman Parkes Research, *Marketers and AI: Navigating New Depths*, 2025.

05 THE FIRST MOVE

None of this requires a transformation programme. It requires four honest decisions, made in order.

01 **Audit for redesign, not addition.**
Before buying anything else, find the one workflow AI could replace a step of — not speed up, replace — and prove it there before scaling anywhere.

02 **Build a judgment gate before you build anything else.**
Someone senior needs to be the last human check on anything AI-touched before it reaches a customer, with real authority to kill it.

03 **Measure outcomes, not activity.**
More content, more drafts, more variants — none of it is a result. Pick one business metric AI is supposed to move, and kill anything that only moves output.

04 **Decide in-house versus partner deliberately.**
Don't let it happen by default because someone senior read a headline. Cost-driven in-housing and capability-driven in-housing are not the same decision.

06 RISK & GOVERNANCE

This isn't a cybersecurity report — there are better people than me to write that. This is the four risks that live specifically inside a creative and marketing operation once generative AI is in the workflow, the ones a general IT security policy won't catch, and what credible bodies currently recommend doing about each.

Shadow AI — the brief in the wrong inbox

LIVE RISK

80% of major brand owners have concerns about their agencies' use of generative AI; 66% cite legal risk as a barrier; only 36% have contractual terms actually governing partner GenAI use. The exposure is briefs, unreleased assets and client data pasted into consumer-grade tools nobody sanctioned.

Mitigation — A named approved-tools list, a one-line rule on what never leaves the building unredacted, and enterprise (not personal) accounts for anything client-facing.

Who owns what you just made — IP & copyright exposure

LIVE RISK

Brand IP counsel across AXA, Diageo, IKEA, L'Oréal, Mars and others name six concrete risks: loss of control of brand data, breach of third-party IP on input, inadvertent infringement in output, no exclusive ownership of AI-made assets, and fraudulent AI ads impersonating the brand. UK case law is actively moving — Getty Images' core copyright claim against Stability AI failed in November 2025, while Disney and Universal's case against Midjourney remains open. Ownership and liability here are not settled law.

Mitigation — Treat AI-output ownership as unresolved until your lawyer says otherwise; keep a record of prompts and sources for anything published; ask vendors (Adobe Firefly is one) whether they offer IP indemnification.

The brand that never said that — impersonation & deepfakes

EMERGING

WPP's own CEO was targeted by a cloned-voice deepfake scam attempting to defraud a colleague. Separately, a majority of consumers report encountering fraudulent celebrity or brand endorsements online, with a meaningful share losing money to them — luxury and consumer-tech brands are the most commonly spoofed.

Mitigation — Adopt a content-provenance standard (C2PA / Content Credentials) for anything official, and give staff a single, quick internal channel to flag a suspected impersonation the moment they see one.

Betting the workflow on one model

REASONED RISK

No public incident yet documents a mid-campaign model update silently changing brand output — this is a plausible operating risk, not (yet) an evidenced one, and it should be argued as such, not dramatised.

Mitigation — Keep at least one alternative tool qualified in every AI-dependent step of the stacks in Chapter 4, so a vendor change is an inconvenience, not a rebuild.

The governance baseline

None of this is exotic. The IPA and ISBA's industry principles for generative AI in advertising — human oversight, mutual AI disclosure, IP due diligence, no undisclosed deepfakes — have been the UK reference framework since 2023. What's changed is the gap between belief and practice: 82% of major advertisers say AI transparency is essential to brand reputation, but only 67% have internal disclosure rules to back it up.

Creative governance, as a decision a team can actually run

A principle is only useful once it's a gate someone checks before publishing. This is the governance baseline above, redrawn as the four questions that decide whether AI-touched work is safe to ship — and what happens on a 'yes' to each.

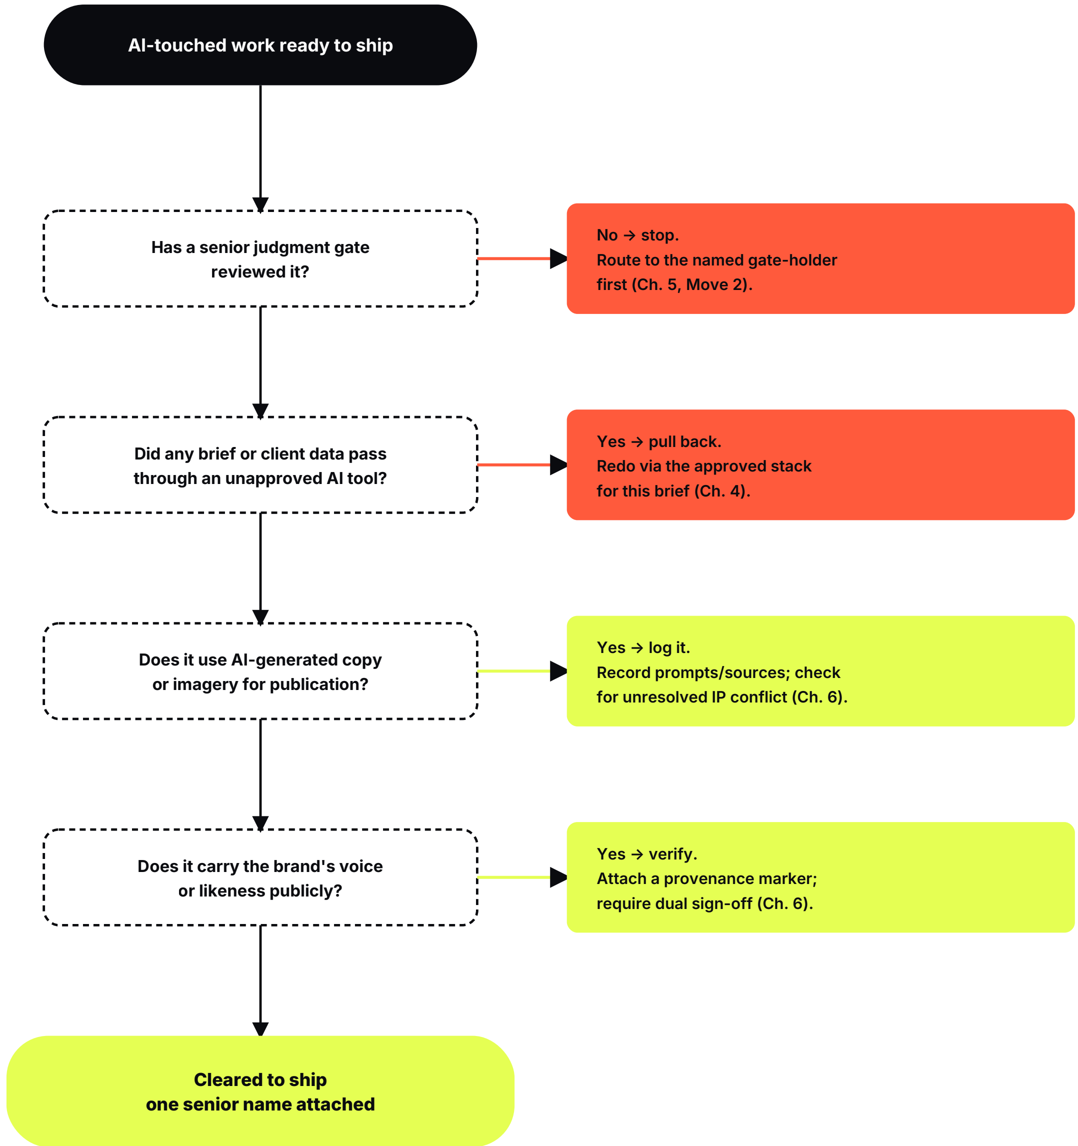

The judgment gate from Chapter 5 isn't just a quality control. It's your risk control too. One senior human, one authority to say no, covers most of this chapter at once.

07 WHERE THIS FITS

None of this is theoretical for me — it's the operating model I already run, across brand strategy, campaign delivery, creative operations and fractional creative leadership.

Clarify the proposition	Brand strategy, positioning, messaging and identity systems — Chapter 4's Stack A.
Launch or grow something	Campaign strategy, digital experience and launch creative — Chapter 4's Stack B.
Fix how work gets delivered	Creative operations, governance and AI-enabled workflows — Chapter 4's Stack C, and Chapter 6 in practice.
Add senior leadership	Fractional creative direction and delivery oversight, without building a full internal team first.

The tools change. The standard, and the judgment behind it, doesn't.

None of this is complicated. It's just uncomfortable to do inside a business that's already moving — which is exactly why most organisations don't do it, and exactly why the ones that do end up with a real advantage over the ones that just bought the tools.

If you want a second pair of eyes on where your operating model actually stands before your next AI decision, that's the conversation I have for a living.

Let's talk →

Travis Hurlock

travis@travishurlock.com · travishurlock.com

CHAPTER 9

THE GOVERNANCE CHECKLIST

Everything in this paper, reduced to the ten questions worth asking before your next AI decision. Print this page.

- [ ] Is there one senior person named as the judgment gate on AI-touched work, with real authority to kill it before it ships?
- [ ] Have you found the one workflow AI should *replace a step of* — not just speed up — and proven it there first?
- [ ] Is there one business metric, not an activity metric, that AI spend is actually accountable to?
- [ ] Was in-house vs. partner a deliberate capability decision, or something that happened by default?
- [ ] Is there a named, approved-tools list, so “shadow AI” has a clear alternative, not just a ban nobody follows?
- [ ] Do briefs and client data have a one-line rule for what never goes into a consumer-grade AI tool unredacted?
- [ ] Is there a record of prompts/sources for anything AI-assisted that gets published, given how unsettled AI-output ownership still is?
- [ ] Does anything “official” carry a provenance marker (C2PA / Content Credentials), and is there a fast internal channel to flag a suspected impersonation?
- [ ] Is at least one alternative tool kept qualified in every AI-dependent step, so a vendor or model change is an inconvenience, not a rebuild?
- [ ] Do your internal AI-disclosure rules actually match what you'd be comfortable telling a customer you do?

CHAPTER 10

ABOUT THE AUTHOR

Travis Hurlock is a UK-based Creative Director and Head of Design, with 30+ years leading brand, UX and creative-operations work — from founder-led startups to a global blue-chip lending division, across organisations including GE, Johnson & Johnson, SAS, cinch, Brightstar and weServed™.

The operating model behind this paper — six senior roles, AI tooling doing the acceleration, one accountable lead holding the standard — isn't theoretical. It's the system currently running client work today, covering brand strategy, campaign delivery, creative operations and fractional creative leadership.

Based an hour from London. UK-wide, remote-first.

weServed™	Brand, UX and creative operations for a UK veterans' support platform — MVP through national launch, including the eight-stage governed production workflow referenced in Chapter 4.
EasySource / Brightstar Financial	Product, UX and creative direction for a fintech case-management platform — 87% of surveyed introducers rated it "made easy."
commonSENSE Monitoring	Brand and product design for a regulated IIoT cold-chain monitoring platform.

If your organisation is somewhere in this paper — spending on AI without the operating model to back it up, unsure where the judgment gate actually sits, or just tired of vendor pitches dressed up as strategy — that's the conversation worth having.

Let's talk →

Travis Hurlock

travis@travishurlock.com · travishurlock.com · LinkedIn

How that conversation usually starts

No deck, no discovery workshop billed by the day. Three steps, in this order:

01

A 30-minute call

Where AI actually sits in your operation today, and where the judgment gate is — or isn't. No pitch, just a straight read.

02

One named brief

Not a full audit. One real piece of work — a brand, a launch, a workflow — scoped against the model in Chapter 4.

03

Fractional or project

Ongoing senior direction, or a fixed-scope delivery. Decided after the first brief, not before it.

APPENDIX A

SOURCES

Every statistic and claim in this paper is attributed below, dated to its original publication — not restated as more recent than it is. Sources checked and current as of September 2026. Where a figure originates from a vendor's own telemetry rather than independent research (marked accordingly in Chapter 6), it is presented as that vendor's claim, not as independently audited fact. Two commonly repeated risk narratives — AI model drift disrupting live campaigns, and vendor lock-in cost — are argued in this paper as reasoned operational risks, because no public, citable incident was found to evidence either; that distinction is kept visible throughout rather than smoothed over for a cleaner story.

1. Gartner, *2026 CMO Spend Survey* — gartner.com
2. MIT NANDA / Initiative on the Digital Economy, July 2025, as reported via deepmarketing.it
3. Klaviyo / Datalily consumer trust study, Dec 2025, with YouGov — reported via emarketer.com
4. Adobe, *The Search for Impact in an Era of Speed*, 2026 — business.adobe.com
5. IPA Agency Census 2026, reported in The Drum — thedrum.com
6. James Hurst, VP Global Creative, HelloFresh, via Superside — superside.com
7. World Federation of Advertisers, *Brands' Concerns on Agency Use of GenAI*, Sept 2024 — wfanet.org
8. World Federation of Advertisers, *Six Key IP Risks of GenAI in Marketing*, Jan 2025 — wfanet.org
9. World Federation of Advertisers, *Global Brands Call for Clearer Consensus on AI Labelling*, April 2026 — wfanet.org
10. *Getty Images v Stability AI*, England & Wales High Court, 4 November 2025 — summarised via lw.com
11. *Disney & Universal v Midjourney*, filed June 2025, ongoing — reported via cnbc.com
12. UK Government, *Report on Copyright and Artificial Intelligence*, 2025/26 — gov.uk
13. Adobe Firefly enterprise IP indemnification terms — helpx.adobe.com
14. AIAAIC incident repository, WPP CEO deepfake scam — aiaaic.org
15. McAfee, *2025 Most Dangerous Celebrity Deepfake Deception List*, Nov 2025 — businesswire.com
16. McAfee, *Most Impersonated Brands of Holiday Season 2025* — via barchart.com
17. IPA / ISBA, *Industry Principles for Generative AI in Advertising*, 2023 (active reference framework) — ipa.co.uk
18. UK Information Commissioner's Office, response to its Generative AI consultation series, Jan 2025 — ico.org.uk
19. IBM, *Cost of a Data Breach Report 2025* — ibm.com
20. Jasper, *State of AI in Marketing 2026* — jasper.ai
21. SAS & Coleman Parkes Research, *Marketers and AI: Navigating New Depths*, 2025 — sas.com

APPENDIX B

THE FINE PRINT

This document is published by Travis Hurlock ("the Author"), trading as Travis Hurlock / travishurlock.com, for general information purposes only.

Not professional advice

Nothing in this document constitutes legal, financial, cybersecurity, regulatory or other professional advice. The Risk & Governance chapter in particular is a general briefing on publicly reported risks and industry commentary, not a security or legal audit of any organisation, and should not be relied upon as a substitute for advice from a qualified solicitor, data protection professional or information security specialist. Readers should seek independent professional advice before acting on any matter referred to in this document.

Accuracy and third-party sources

Statistics, quotations and claims in this document are drawn from third-party research, reporting and publicly available sources current as of the publication date and cited in Appendix A. The Author has made reasonable efforts to cite accurately but does not warrant the completeness, accuracy or continued currency of any third-party statistic, study or claim, and figures may have been superseded by later research. Legal matters described (including ongoing litigation) reflect the position at the time of writing and may have changed.

No warranty; limitation of liability

This document is provided "as is" without warranty of any kind, express or implied. To the fullest extent permitted by law, the Author accepts no liability for any loss or damage arising from reliance on its contents.

Trademarks

Product and company names mentioned in this document — including but not limited to Claude, OpenAI, Gemini, Midjourney, Runway, ElevenLabs, Figma, Adobe Creative Cloud, Adobe Firefly, Notion, Asana, Zapier, HubSpot, Mailchimp, Google Analytics, Semrush, SimilarWeb, Metricool, Gartner, MIT, Adobe, IPA, ISBA and WFA — are trademarks or registered trademarks of their respective owners. Their use here is for identification and reference purposes only and does not imply endorsement of, or affiliation with, the Author.

Copyright