

Matt Kundo

Austin, TX · (318) 470-4223 · mattkundo@mattkundodigitalmarketing.com · linkedin.com/in/mattkundo · texascommercialplans.com

GTM Engineer · Marketing operations, attribution and AI systems

I'm a marketer who builds the internal tools. Ten years across paid media, SEO, analytics and CRM operations, and since December 2024 I've spent most of my time engineering the measurement and automation rather than doing it by hand. I've run an agent platform in production since January 2026, rebuilt in May 2026 on the Claude Agent SDK, executing marketing operations across a five-property book with every irreversible action behind a switch that fails closed. I ran all of it manually first, which is why the automation covers the loops that matter. I'm winding the agency down to build one system properly instead of five in parallel.

MEASUREMENT AND ATTRIBUTION

- **Built the unified attribution framework** that first separated a lead from a completed enrollment for a retail electricity provider, across paid, direct mail, call tracking with disposition tagging, and offline enrollment records matched back to touchpoints. It's what made cost per closed customer measurable instead of cost per lead.
- **Found two Google Ads conversion actions that had recorded zero conversions for their entire lifetime.** The enrollment records had never been connected to an import, so bidding was optimizing toward a signal that didn't exist. I built the uploader on the Ads API using Enhanced Conversions for Leads, hashing email and phone before transmission, and set it observation-only on the first load so a backfill couldn't disturb live bidding: 194 of 197 accepted at 98.5%, \$48,570 in conversion value. It's run daily since.
- **A second uploader that fails closed on purpose.** The zip-code lead form attracts spam, so an automatic uploader would have taught bidding to chase junk. This one pushes only click IDs a human has marked legitimate: 2,836 rows scanned, 56 marked, 2 skipped as fraudulent.
- **Killed a reporting request instead of delivering it.** Asked for revenue by promo code, I found the parameter was event scoped, so it read as not set across all 17,181 key events that year. Sessions by promo code worked fine, which is why nobody had caught it. No fix to the tag would have produced the report I was asked for, so I specified the replacement and said plainly I'd found the reason it was impossible, not that I'd measured a lift.
- **Channel-level attribution in Tag Manager**, as four custom JavaScript variables mapping 1,107 promo codes into GA4, reduced from 1,177 raw by an audit that found 69 duplicates and one code double firing across two segments. **Customer Match** on Google's Data Manager API after the Ads API path was blocked, first ingest 28,032 members.

PLATFORM AND AI SYSTEMS

Multi-agent marketing platform (Python 3.13, Claude Agent SDK, MCP, Postgres, pgvector) · in production since Jan 2026, current version since May 2026 · built and operated solo, against real client accounts and real spend

- **Agents, tools, and an MCP server.** 29 agent roles, 19 of them carrying production run history in the last 30 days, coordinating through structured JSON handoffs in Postgres rather than free-form chat, with a hard two-hop limit per task. An in-process MCP server exposes 390 agent-callable tools. I scope permissions inside the tools themselves, because the SDK's allowlist auto-approves rather than restricts and isn't a boundary.
- **Guardrails I'd want if it were my own brand.** Every irreversible or outward-facing action sits behind one of 44 database-backed kill switches that fail closed, so an unseeded flag denies the action rather than allowing it, and a deterministic compliance gate has hard blocked 148 outputs in production. A completion gate stops an agent marking work done unless it self-certifies with checkable evidence, which I built after learning that agents report success most confidently when no tool exists for the thing they were asked to do.
- **A content pipeline with evals that block.** An article runs research, drafting, quality review, image sourcing, CMS document, publish and indexing, with an information-gain floor scoring each draft against the live top ten and an LLM judge whose verdict is allow-listed to one exact string. A non-pass naming no concrete fix is treated as incoherent rather than trusted, because a judge that can only say "worse" teaches you nothing.
- **Cost and correctness.** Over 110 scheduled jobs run in production and roughly a third never call a model at all. That split is the cost control and a quality decision too, because a scheduled writer asked what should happen next always answers, and always answering is how a system manufactures work nobody needed. Over 6,700 tests, green in about 75 seconds, on CI for every pull request.

DATA AND PRODUCT

- **A pipeline that was reporting deletions it had invented.** A price-comparison product I own runs an extraction pipeline into Supabase, and its weekly report used to lead on how many plans had left the catalog. That number was mostly the extractor missing rows. Across 30 consecutive scheduled runs it returned between 11 and 36 percent of the active catalog each time, while the lifecycle rule retired anything missed twice, so a plan it had simply failed to find was gone in about a day and a half. The root cause was in the upsert key, which carried a model-generated plan name, so a rename orphaned the old row and minted a new one. I tested the obvious fix and killed it: the premium extractor tier returned fewer plans, renamed them so no keys matched, and ran two to eight times slower.
- **Shipped texascommercialplans.com end to end and solo.** 13,280 plan records, 8,177 of them active,, from a 148-row provider registry of which 29 are active, refreshed twice daily, plus Stripe checkout with idempotent webhook handling across 10 Supabase Edge Functions. Postgres and Supabase with applied migrations and pgvector semantic memory behind a swappable embedding backend.

OUTCOMES

Client work, anonymized by agreement.

- **101% ROI on a \$76,279 channel mix**, October 2025: \$153,610 revenue, 594 enrollments at a blended cost of \$128.45 each. An enrollment is a closed customer rather than a lead, so that's a cost per closed deal, and it came out of the attribution framework above rather than platform-reported conversions.
- **Launched paid search from zero** for an at-home health testing e-commerce brand. There was no paid channel the month before launch. Cost per conversion fell from \$253 in the launch month to \$38 six months later, and monthly conversions grew from single digits to 39.
- **42.6% conversion rate** on the stronger of two rebuilt paid landing page variants, against 23.3% on the other.

EXPERIENCE

Matt Kundo Digital Marketing, Austin, TX · Founder · Jul 2018 to Present

Ran a portfolio of client and owned properties solo with direct P&L responsibility: paid media, SEO, content, analytics engineering, attribution, CRM operations and web development. Built and now operate the platform above, which runs delivery for that book and is the reason one person can carry it.

Rebel Interactive Group · Marketing Strategist, SEO and Media Specialist · Feb 2017 to Jul 2018 · **FlexTekk** · Marketing and Sales Manager · Sep 2015 to Jan 2017 · **Lucky Robot, Austin, TX** · Marketing Coordinator · Jul 2013 to Jul 2015

TECHNICAL SKILLS

Languages Python 3.13, SQL, JavaScript

AI and agents Claude Agent SDK and Claude Code as a daily driver, Model Context Protocol including an in-process MCP server, multi-step LLM pipelines with structured outputs, LLM-as-judge evaluation, RAG on pgvector

Ad platform and measurement APIs Google Ads API including Enhanced Conversions for Leads and conversion-action management, Google Data Manager API, GA4 and its Data API, Google Tag Manager API and custom JavaScript, Search Console, Meta Ads, CallTrackingMetrics, DataForSEO

Data and infrastructure Postgres with migrations, Supabase, Azure SQL, pgvector, Python data pipelines, webhooks and idempotent event handling, OAuth and token lifecycle, Cloudflare, Netlify

Revenue systems HubSpot at the API level including lead-source attribution, pipeline inventory and win/loss reporting, Stripe, QuickBooks API, Power BI

Practice pytest, git, GitHub Actions, CI on every pull request

EDUCATION

The University of Texas at Austin · B.A. International Relations and Global Studies · Dec 2014 · McCombs Business Foundations Certificate

Every number on this page traces to a system of record, and I'm happy to walk through any of them.