

Custom Marketing Attribution in 6 Weeks — Without Touching Salesforce

Pallas Advisory | Marketing Attribution Case Study

The Situation

A global, multi-brand B2C company was spending over a million dollars annually on paid media across Google, Bing, and Meta — and couldn't tell you which campaigns actually drove revenue.

The marketing team relied on manual Excel reconciliation. Every week, the analytics team spent 15+ hours exporting sales reports from Salesforce, pulling spend data from ad platforms, and trying to match them in spreadsheets. Results arrived three weeks after campaigns ran. By the time anyone saw whether something was working, the budget had already been spent.

The real problem was deeper than slow reporting. Of the 34 tracking parameters that ad platforms send to landing pages, only 7 survived into the company's lead processing system. The other 27 — campaign IDs, ad group data, keyword matches — were lost at the handoff between the landing page platform and the CRM. By the time a lead became a sale, attribution was guesswork.

Their external media buying agency provided performance reports, but the company had no independent way to verify them. The agency used proprietary tracking parameters on their own systems. The marketing team was making multi-million-dollar budget decisions based on numbers they couldn't validate.



Why They Didn't Just Buy Something

This wasn't a case of not knowing vendor solutions existed. They evaluated the options and rejected all of them for specific reasons.

Salesforce-native attribution would have required custom objects, Apex triggers, and Marketing Cloud integration. I've built custom solutions inside Salesforce before — I know exactly what that level of effort looks like. The problem: the company's technology infrastructure had been through years of mergers and acquisitions, and the systems hadn't evolved to keep pace. All custom systems ran on legacy code, and the IT team — stretched thin maintaining what already existed — didn't have bandwidth for new projects. Any Salesforce customization meant navigating politics across three departments, with a realistic timeline of 6+ months and a budget north of \$200K. Worse, the data loss was happening before Salesforce — building inside the CRM would have meant working with already-degraded data. I felt strongly enough about this that if the direction had been to build inside Salesforce, I would have recommended they find someone else for the project. That's how confident I was that there was a better way.

Enterprise SaaS platforms like Marketo Measure came with \$70K+ annual licensing — more than the entire cost of a custom build, every single year, with the data living on the vendor's infrastructure.

The external agency pushed its own solution — a third-party platform running on the agency's instance. This would have created maximum lock-in: the company's attribution data living on their vendor's platform, making it impossible to independently validate the vendor's own performance claims. That's not attribution. That's trust.

A cloud data warehouse (Snowflake, BigQuery) was considered and rejected as overkill — roughly \$2,000/month for a use case that needed \$50/month of infrastructure.

The core requirement was simple and non-negotiable: the company needed to own its data, on its own infrastructure, in a standard format it could migrate anywhere.



What I Built

The goal was simple: give leadership a single dashboard where they could see return on ad spend by campaign, by brand, by country, and by time period — updated daily instead of three weeks later.

I designed and implemented a custom attribution system that captures marketing touchpoints at the source — before any data is lost — and connects them to actual revenue outcomes. The system pulls in ad spend from Google, Bing, and Meta automatically, matches form submissions to closed sales, and produces the ROAS reporting that had been impossible with the previous setup.

The system was built with GDPR compliance baked in — row-level security, audit trails, PII handling designed for EU data requirements — and designed for handoff. The analytics team now operates it independently, with daily monitoring requiring about 5 minutes and weekly maintenance taking 15-20 minutes. They were told to escalate to me only if the core processing functions break — everything else, they own.

The most important design decision was data ownership. The company's data lives in a standard database that they own and control. Not on a vendor's platform. Not inside Salesforce. Not locked into a SaaS subscription they can't leave. If they ever want to switch providers, migrate to a different infrastructure, or bring the whole thing in-house — they can. Try doing that with a \$70K/year attribution platform.

What the System Uncovered

Building a custom attribution system doesn't just produce dashboards. It produces truth. And sometimes the truth is uncomfortable.

A \$1.28M currency discrepancy caught within weeks of deployment. In multi-currency organizations, the same sales data files can contain both local currency and USD amounts with ambiguous labeling. Any new system consuming that data will hit the same trap — and ours did. During validation, I identified that the system was misinterpreting foreign currency amounts as USD, creating a \$1.28M discrepancy that would have silently corrupted every metric the system produced going forward. Building custom attribution forces you to understand every field in the data, not just connect APIs. That validation process is what caught it. An off-the-shelf tool that simply ingests data without questioning it wouldn't have.

A quarter-million-dollar spend understatement. Ad platforms report spend broken out by device and network type. The initial data load was capturing subsets instead of summing all rows, understating 23% of total campaign spend. This was



discovered during a training session with the analytics team — they noticed the numbers looked low compared to what they saw in the ad platforms directly. Every ROAS figure reported before the fix was artificially inflated. The fact that the team caught it in real time, and watched the diagnostic process, actually built more trust than if everything had worked perfectly from the start.

ROAS varied dramatically across product lines — and nobody knew. The system revealed that some product lines were generating over 2x return on ad spend while others were barely breaking even. That's the kind of insight that changes how you allocate a multi-million-dollar media budget. It's also the kind of insight that's invisible when you're relying on agency reports and Excel reconciliation.

The Numbers

Metric	Before	After
Time to campaign performance data	3 weeks	Next day
Weekly manual reconciliation work	15+ hours	Eliminated
Tracking parameters captured	7 of 34	34 of 34
Marketing touchpoints in system	None (spreadsheets)	50,800+ captured and counting, 44,400+ backfilled
Cross-platform ROAS visibility	Nonexistent	By campaign, brand, country, time period
Data ownership	Agency-dependent	Company-owned, fully portable
Annual infrastructure cost	N/A	Under \$1,000/year

Timeline: 4 weeks from build start to first live data in dashboard. 6 weeks to executive-facing ROAS reporting. The analytics team was trained and operating independently within 8 weeks.



What Stakeholders Said

The internal project sponsor, who had visibility across the entire organization, described the engagement this way:

"I wake up every morning and there's three or four emails that are extremely professional, extremely thorough, detailed, super helpful. I am seeing two people who are working extremely well together, extremely efficient and extremely productive."

The CEO's first reaction to the live dashboard was to immediately start exploring filters by brand and by country — then articulate exactly the question the system was built to answer: "Within a brand, within a country, what campaigns are working, what are not, what should I try to buy more of? Because that's what we have to tell the media buying group."

The newly hired marketing analyst, seeing the dashboard for the first time: "I can look at it and tell exactly what you are probably trying to get to."

The analytics team engaged immediately — requesting enhancements, asking informed questions, and running ahead of the training schedule. They went from consumers of stale spreadsheets to operators of a live system.

What Was Hard

I don't believe in case studies that pretend everything went smoothly. Here's what was genuinely difficult.

The ROAS calculation was wrong four times before it was right. The initial figure was 3.77x. It dropped to 1.90x, then 1.66x, before settling into the correct segmented view. Each correction required rebuilding SQL views and revalidating. The root cause was an architectural mistake: spend and attribution were entangled in a single database view, so filtering by attribution model inadvertently filtered out spend data. The fix — treating spend as a fact and attribution as an interpretation, in separate views — became the most important design decision in the entire system. If this had gone to the CEO uncorrected, credibility would have been destroyed.

A core library didn't work in the runtime environment. The original approach for parsing sales data relied on a JavaScript library that couldn't run in the serverless environment. This required a full-day pivot to an alternative approach mid-build. A good reminder to test library compatibility before committing to an architecture, not after.



The match rate looked broken for weeks. The system showed a 10% attribution match rate — far below expectations. Investigation revealed it was a labeling bug: over 1,000 sales were actually matched but still flagged as "unmatched" due to a sequencing issue in the identity resolution pipeline. Real match rate was 42%, with higher rates for data captured after the webhook system went live. The lesson: when a metric looks wrong, check the measurement before questioning the methodology.

Why This Matters Beyond One Company

This project wasn't unusual in its complexity. Most mid-market B2C companies with legacy CRM infrastructure face the same problem: marketing spends real money, but the systems that track that spend were designed for lead management, not attribution. The data loss happens silently, in the handoff between platforms, and nobody notices because you can't miss what you never had.

The traditional fix — a six-figure Salesforce overhaul or a \$70K/year enterprise platform — prices out most mid-market companies. The alternative most settle for — trusting their agency's self-reported numbers — isn't really attribution at all.

What I build represents a new path — one that wasn't possible even a few years ago. Modern infrastructure (managed databases, serverless computing, automated data connectors) means custom systems can now be built faster, cheaper, and more flexibly than legacy enterprise platforms. And because the data lives in open, standard formats the company controls, the system is ready for whatever comes next — without waiting for a vendor's product roadmap or paying for their next feature tier. No vendor lock-in, no perpetual licensing, no IT dependency. A solo consultant who builds working systems, not a consulting firm that delivers strategy decks.

The system described in this case study is now in production, with Phase 2 planning underway to expand channel coverage.

If your marketing team can't answer "which spend actually drove revenue last quarter," I can probably tell you why — and fix it. Get in touch at pallasadvisory.com.

Pallas Advisory builds marketing data infrastructure for mid-market companies — connecting ad spend to actual revenue without six-figure platform costs or year-long IT projects. Founded by Melissa Kolbe, who combines 15+ years of enterprise marketing operations with hands-on technical implementation. Based in the Netherlands. Learn more at pallasadvisory.com.