

Real Time AI Marketing Analytics: +27% ROI For A Global Retail Analytics Company

A unified Medallion data platform, predictive modeling, and automated model drift detection – cutting data prep from weeks to hours and forecasting campaign outcomes before launch. Wondering how AI marketing analytics could reshape your marketing ROI?

Start the conversation with Teamvoy

SERVICES

[AI/ML Engineering](#), [Data Platform Architecture](#), [Medallion Lakehouse Design](#), [MLOps & Drift Monitoring](#), [CRM/ERP Data Integration](#)

INDUSTRY

[Real Time Marketing](#), [Retail Analytics](#)

CLIENT

A global real-time marketing and retail analytics company

MARKETING ROI LIFT:

+27% Achieved within the first six months

DATA PREP TIME REDUCTION ACROSS CRM AND ERP:

60% Manual data preparation dropped from weeks to hours

Before Launch: Building an AI Marketing Analytics Platform

Forecasting Campaigns Before Launch: Building an AI Marketing Analytics Platform

Executive Summary

How Did A Global Retail Analytics Company Turn Fragmented Real Time Marketing Data Into A Predictive ROI Engine?

A global real time marketing and retail analytics company was sitting on the right data and the wrong workflow. Product, campaign, and customer information lived across CRM and ERP systems. Marketing teams needed predictive insight to optimize campaign performance, but every analysis started with weeks of manual data preparation, and once models were finally in production, there was no system in place to detect drift or trigger retraining. The result: insights arrived late, models decayed quietly, and budget decisions ran on dashboards instead of forecasts. This case study walks through how Teamvoy built an AI marketing analytics platform on top of a production-grade Medallion Architecture, integrated it with the client's CRM and ERP, and shipped ML attribution models with automated model-drift detection and retraining pipelines. The result: a 27% lift in overall marketing ROI within six months, manual data preparation cut from weeks to hours, and campaign outcomes that can be forecast before launch, not interpreted after the fact.

About The Client

Who Is The Client And Why Did Real Time Marketing Modeling Need To Become Predictive?

The client is a global marketing and retail analytics company that integrates product, campaign, and customer data across enterprise CRM and ERP systems. Its business depends on translating that integrated data into decisions – which campaigns to scale, which to pause, which budget reallocations actually move the needle.

That dependency reset the team's ambitions. The next phase was not better dashboards; it was forecasting. The client wanted to boost campaign performance prediction and marketing ROI using AI-powered data analytics and automation, and it needed an intelligent analytics platform that could predict marketing campaign outcomes based on product attributes and digital assets – and integrate cleanly with the CRM and ERP infrastructure already in production. Anything that could not plug into the existing enterprise data flow was a non-starter.

The Challenge

What Problem Does AI Marketing Analytics Solve When Data Is Fragmented And Models Decay Silently?

Four challenges defined the baseline, each one easy to describe on a slide, each one expensive to live with at scale.

Marketing and product data were fragmented across different systems. Campaign performance, customer behavior, and product attributes sat in separate stores, with no unified surface a model could be trained against. Every modeling effort started by re-stitching the same data.

There were no predictive insights for optimizing campaign performance. Marketing teams could see what had already happened, but had no model-driven view of what was likely to happen next. Without predictive analytics marketing decisions ran on intuition and recency, not on forecasted ROI. Complex ETL processes slowed down model training and experimentation. Pipelines were brittle and serial, which meant every new modeling idea had to wait its turn behind data engineering work. The cost of trying something was high enough that fewer things got tried.

No proactive system existed for monitoring model drift and retraining. Models went to production and then aged silently. By the time anyone noticed performance was off, the wrong budgets had already been spent, and “retrain it” was a manual project, not a pipeline.

REAL TIME AI MARKETING ANALYTICS · THE CHALLENGE

The right data — and the wrong workflow

Easy to describe on a slide · expensive to live with at scale · global retail analytics baseline

- CHALLENGE 01**
Data fragmented across systems
Campaign performance, customer behavior, and product attributes sat in separate stores — no unified surface a model could be trained against.
Every modeling effort started by re-stitching the same data
- CHALLENGE 02**
No predictive insights
Marketing teams could see what had already happened — but had no model-driven view of what was likely to happen next.
Decisions ran on intuition and recency — not forecasted ROI
- CHALLENGE 03**
ETL slowed every experiment
Pipelines were brittle and serial — every new modeling idea had to wait its turn behind data engineering work.
The cost of trying something was high enough that fewer things got tried
- CHALLENGE 04**
Models decayed silently
No proactive system for monitoring model drift or retraining. Models went to production and then aged quietly.
By the time anyone noticed → the wrong budgets had already been spent

The combined result — insights arrived late, models decayed quietly, and budget decisions ran on dashboards instead of forecasts

Each challenge predictable · each one expensive to live with quarter after quarter

What We Did

How Did Teamvoy Build The AI Marketing Analytics Platform?

The engagement delivered four interlocking workstreams: a unified data platform, predictive ML models, an MLOps layer for drift and retraining, and the integration surface back into the client’s CRM and ERP.

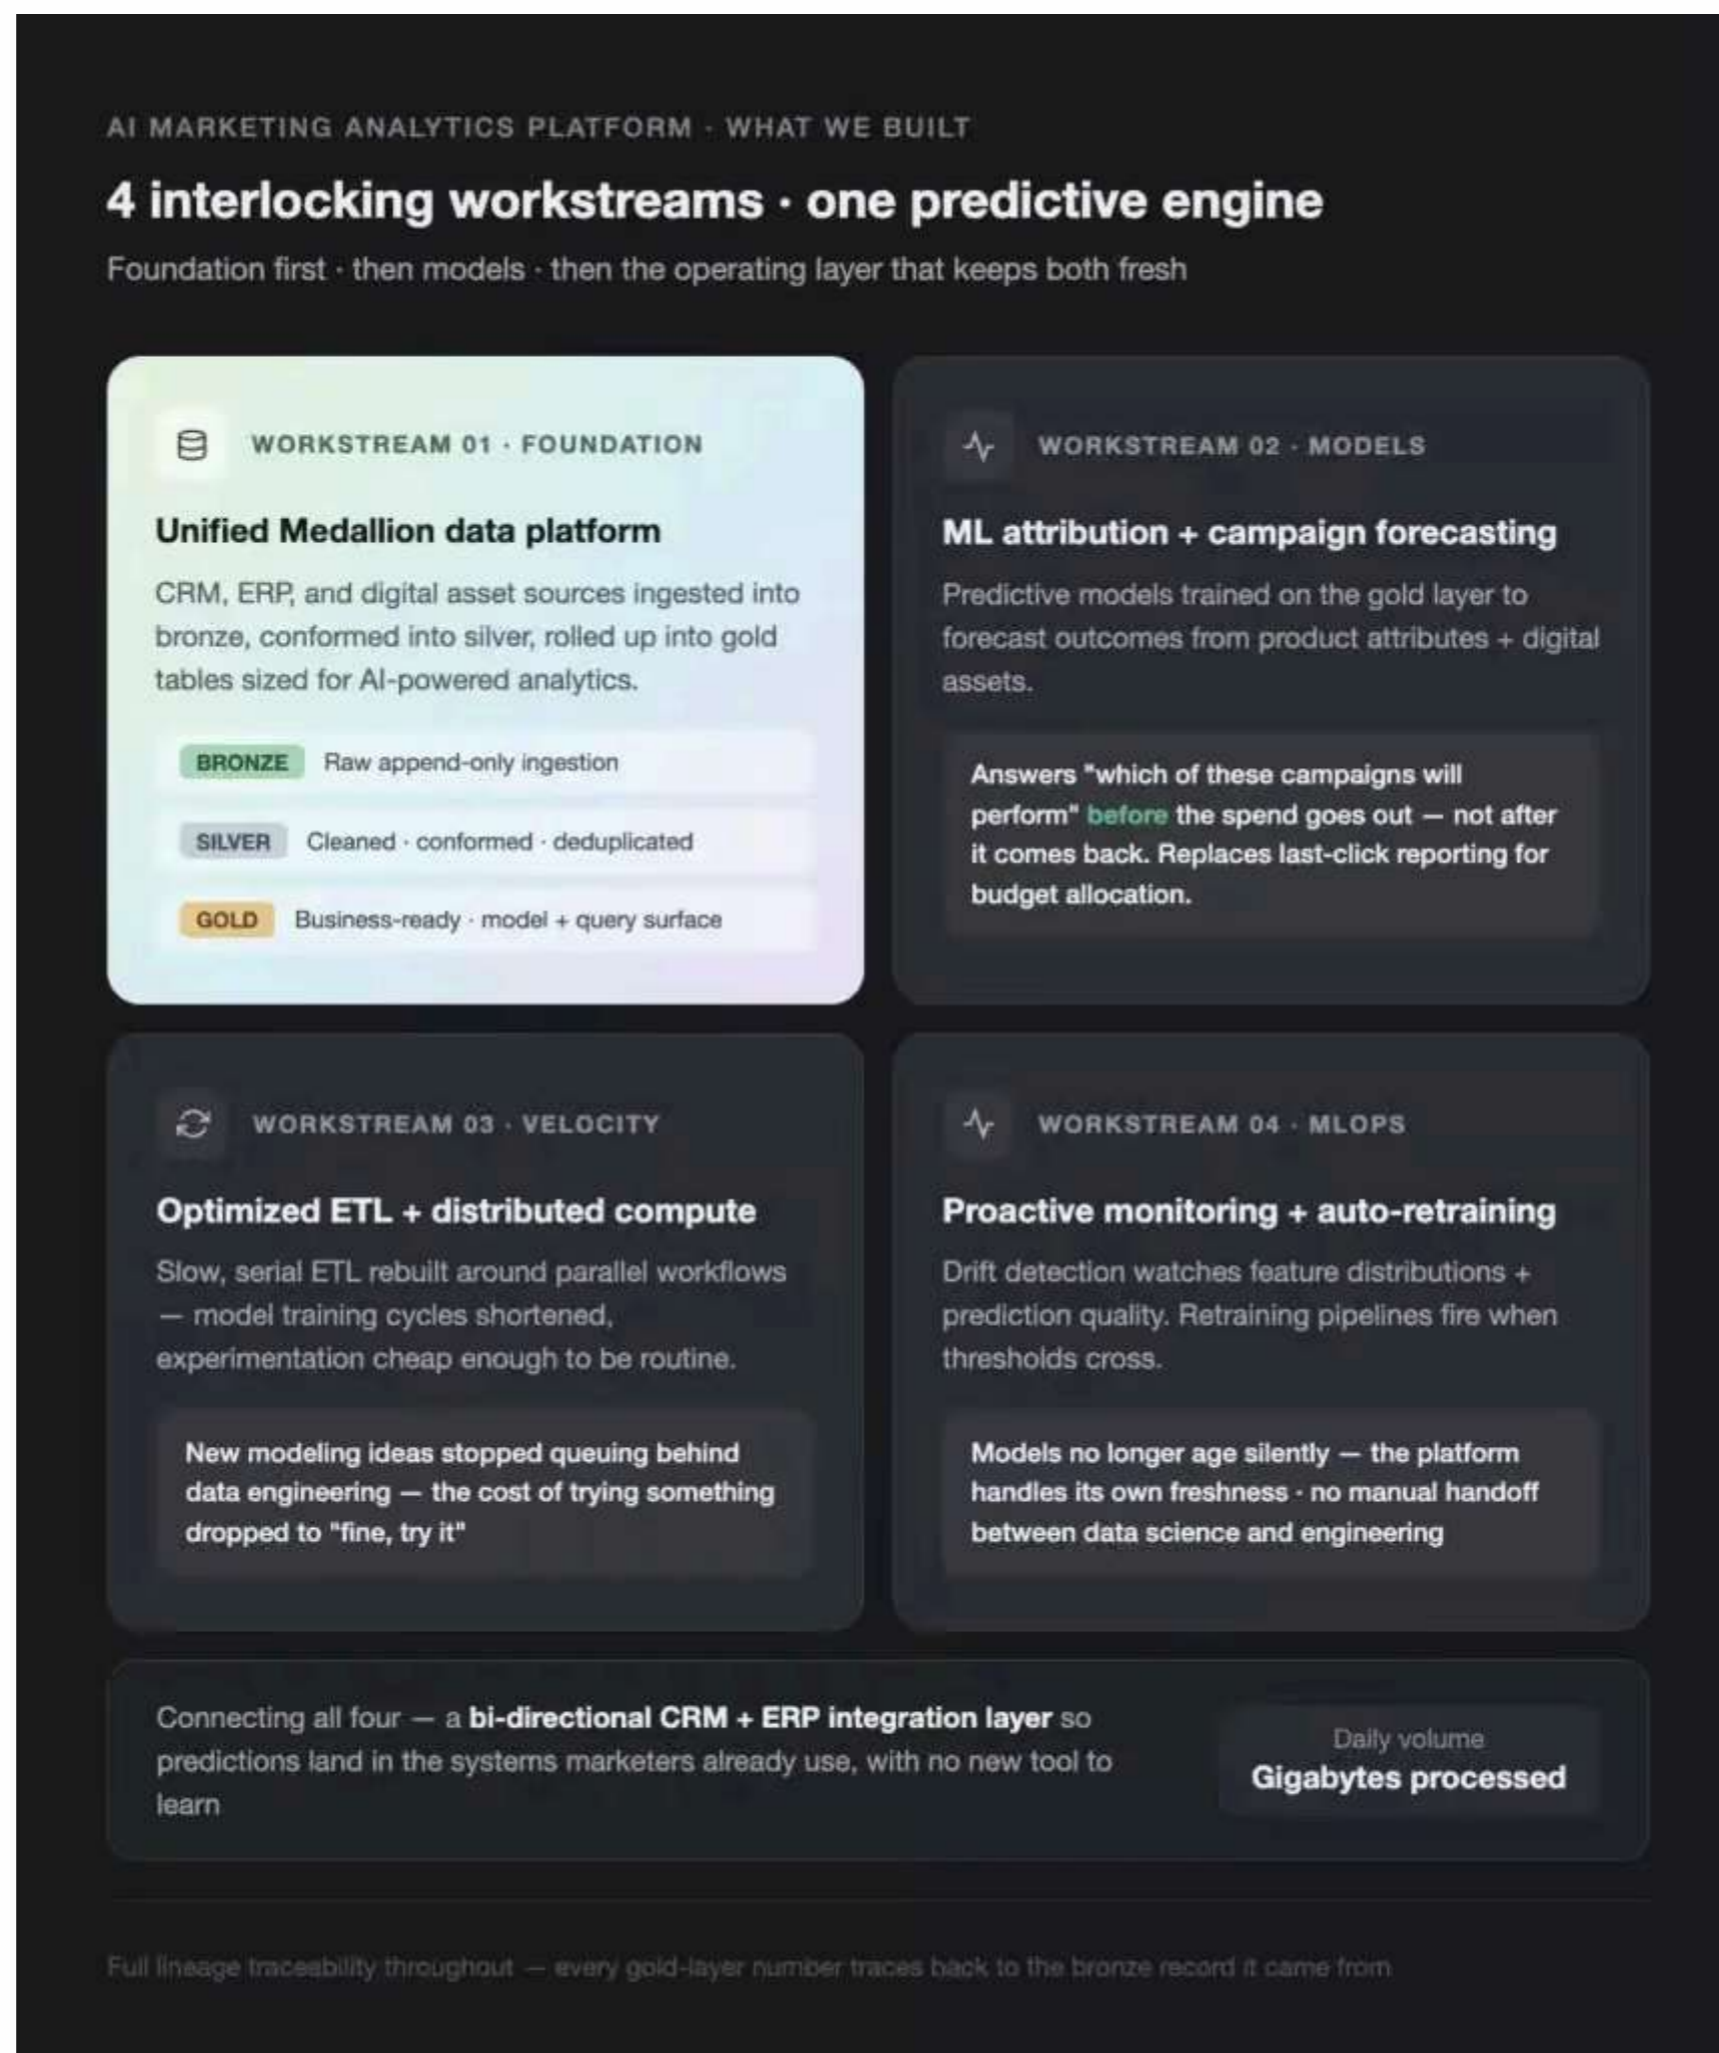
Unified Medallion data platform. Product, campaign, and customer data were ingested from CRM, ERP, and digital asset sources into a bronze layer, conformed into a silver layer, and rolled up into gold tables sized for AI-powered data analytics. The architecture was designed to process gigabytes of daily data with full lineage traceability, every metric a model trained on, or a marketer queried, could be traced back to its source row.

ML-driven attribution and campaign forecasting. Predictive models were trained on the gold layer to forecast campaign outcomes based on product attributes and digital assets, and ML-driven attribution models replaced last-click reporting as the basis for budget allocation. This predictive modeling could finally answer “which of these campaigns will perform” before the spend went out, not after it came back.

Optimized ETL and distributed compute. The slow, serial ETL was rebuilt around optimized workflows and distributed computing, which materially shortened model training cycles and made experimentation cheap enough to be routine. New modeling ideas stopped queuing behind data engineering.

Proactive monitoring and automated retraining. An MLOps layer was built on top of the modeling pipeline, with automated model drift detection that watches feature distributions and prediction quality in production, and retraining pipelines that fire when drift crosses defined thresholds. Models no longer age silently,

the platform handles its own freshness



Tech Stack

Which Technologies Power The Real Time AI Marketing Analytics Platform?

- Medallion Lakehouse Architecture (bronze/silver/gold). Production-grade data layering for gigabytes of daily marketing, campaign, and product data.
- Distributed computing framework, for parallelized ETL and ML training across the platform.
- ML attribution and forecasting models, predictive modelling for campaign outcome prediction based on product attributes and digital assets.
- MLOps pipeline, automated model drift detection, model performance monitoring, and retraining triggers.
- CRM and ERP integration layer, bi-directional flow so predictions land in the systems marketers already use.
- Data lineage tracking, full traceability from gold-layer metrics back to bronze-layer source records.

Key Features

Which Features Define A Production-Ready AI Powered Analytics Platform?

- Pre-launch campaign forecasting, predict outcomes before budget is committed, not after it is spent.
- ML-driven attribution that replaces last-click logic with a model of how channels actually contribute.
- Unified marketing data surface, product, campaign, and customer data conformed in one place, queried as one source of truth.
- Automated model drift detection on features and predictions, with thresholds tuned to the business cost of staleness.
- Retraining pipelines that fire automatically when drift triggers, with no manual handoff between data science and engineering.
- CRM/ERP write-back so insights surface in the workflows marketing teams already operate in.
- Full lineage traceability, every prediction explains itself back to the source data that produced it.

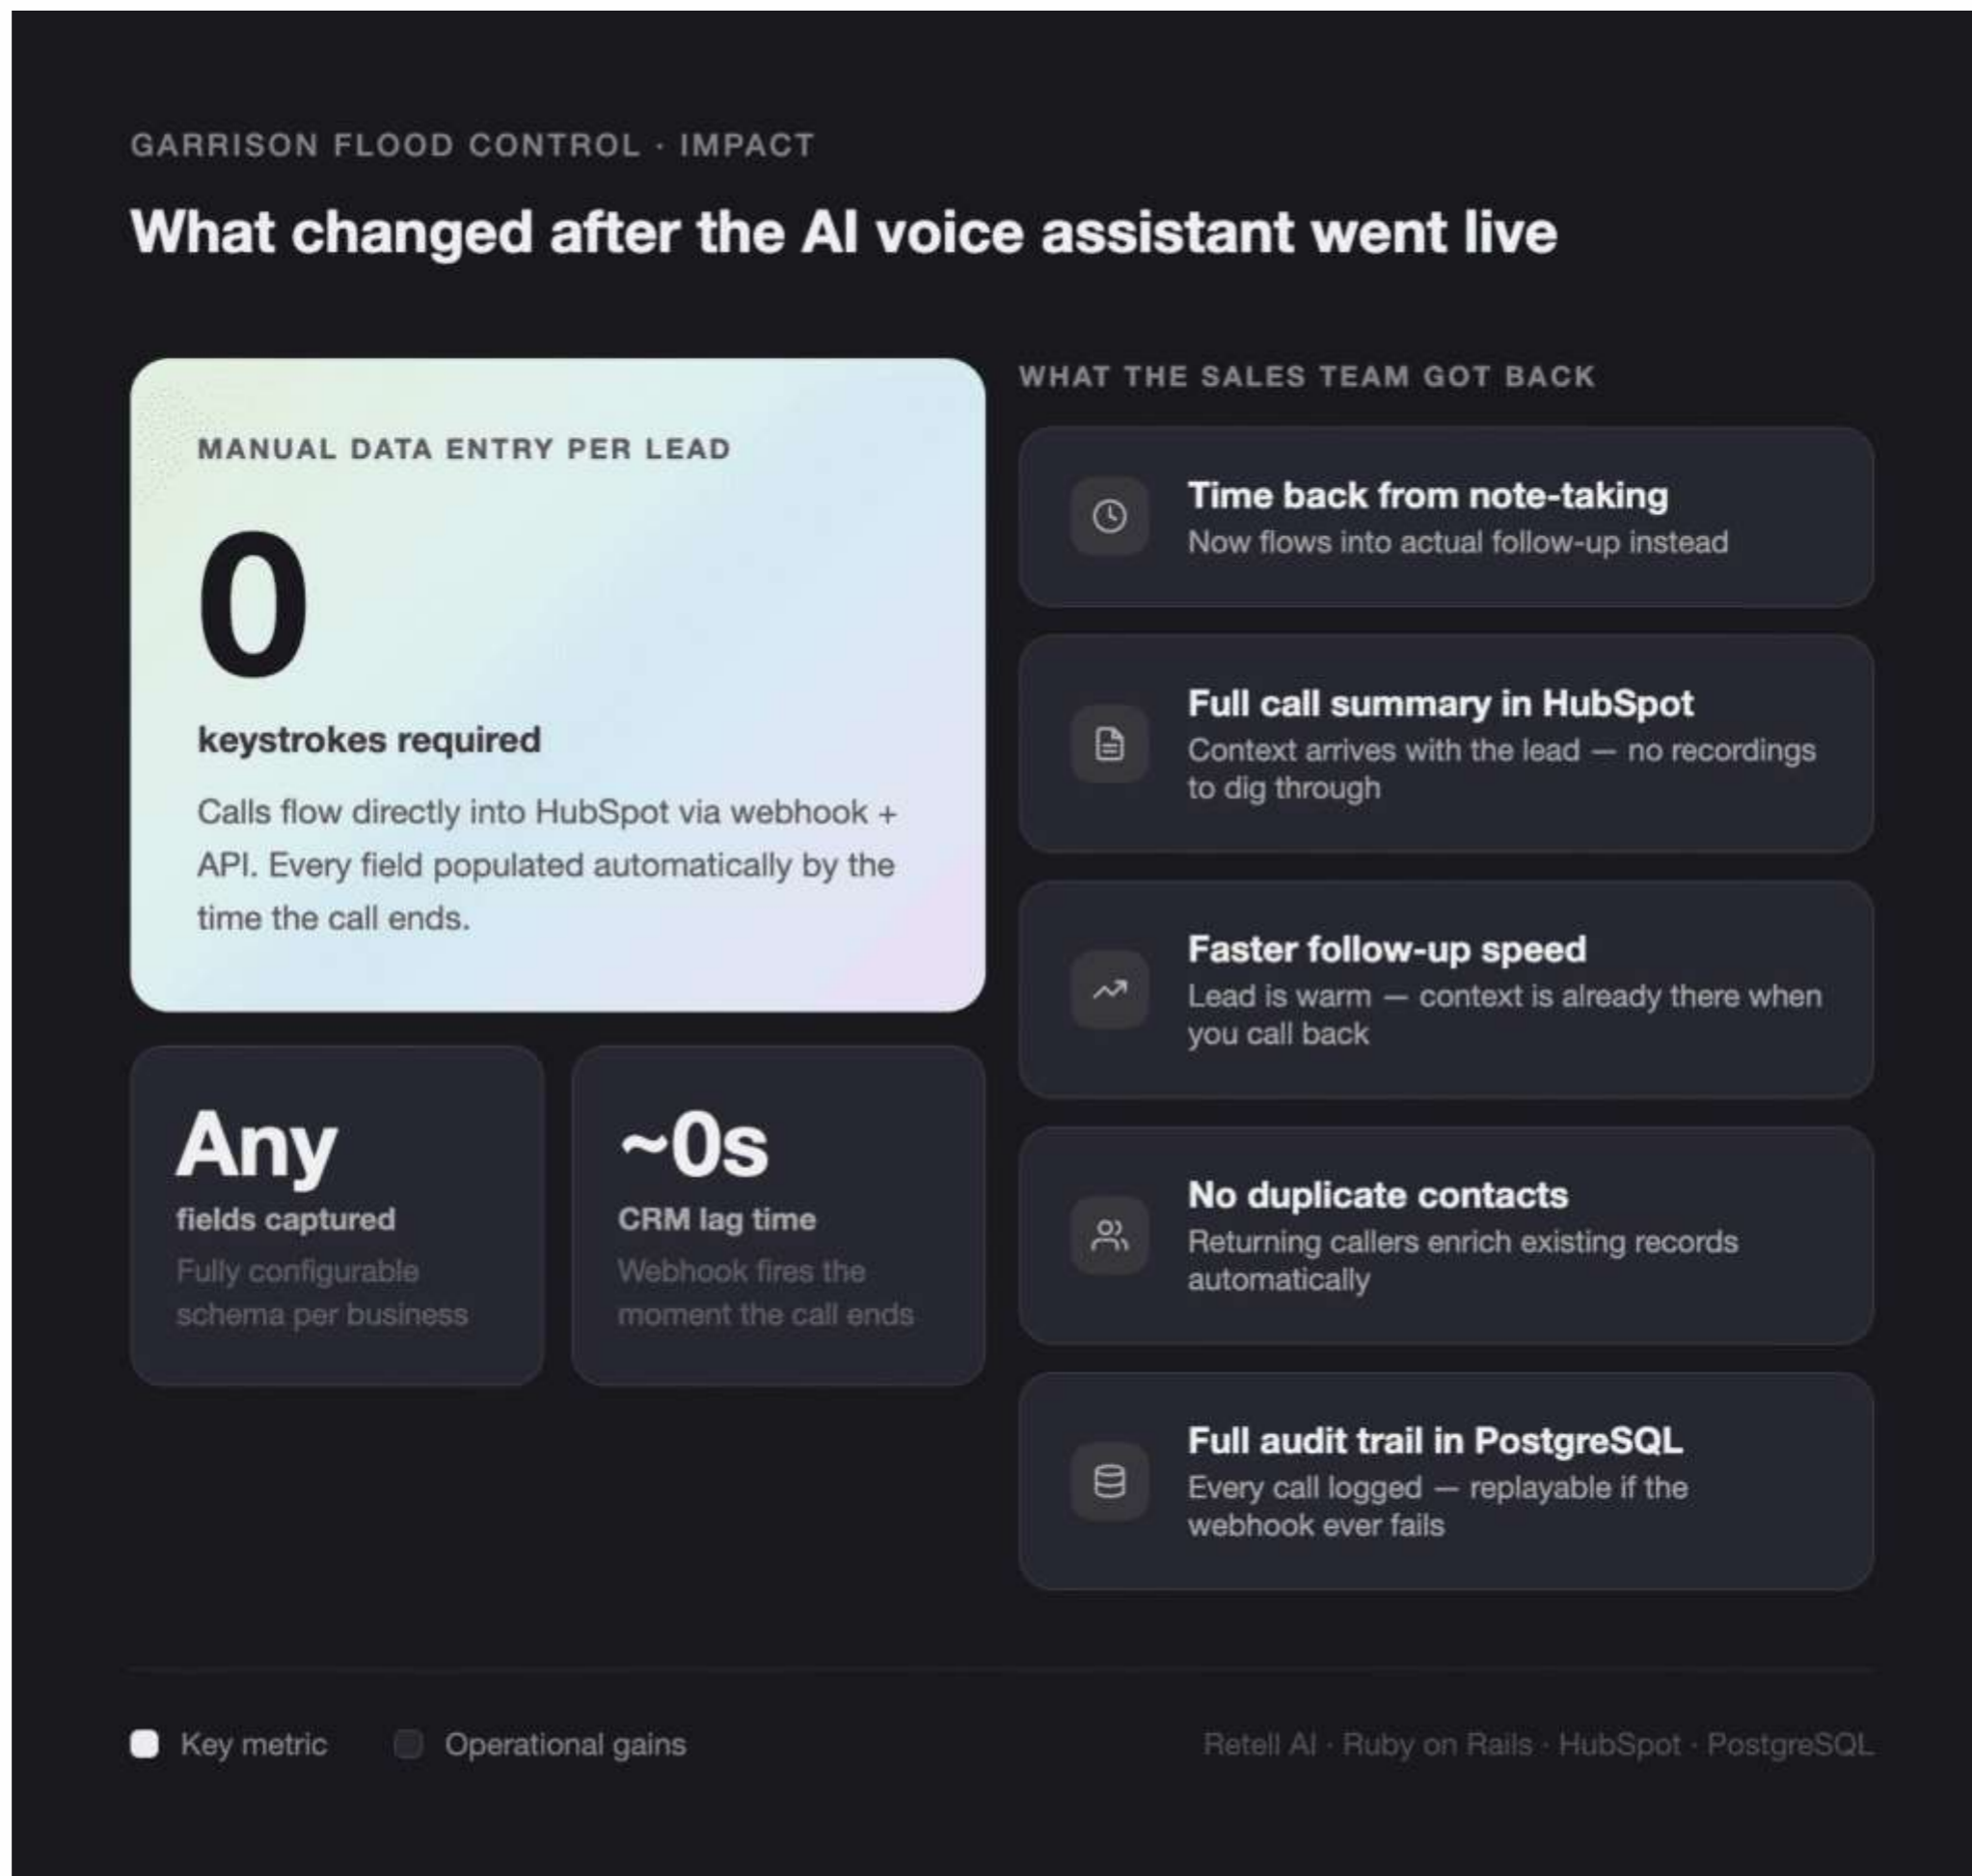
Key Engineering Decisions

Which Engineering Decisions Made The AI Powered Data Analytics Solutions Reliable In Production?

Two decisions shaped how the platform behaved under real marketing load.

Medallion before models. The data layer was built before the modeling work began. The Medallion structure was what made every later round of predictive modelling reproducible, and reproducibility is the property that separates a production ML platform from a notebook.

Lineage as a first-class requirement. Every gold-layer row carries traceability back to its bronze sources. Lineage shortened debugging, made audits trivial, and gave marketing stakeholders enough confidence in the numbers to actually change budget allocations on the strength of them.



Impact

What Impact Did AI Marketing Analytics Have On The Business?

The platform's measured impact moved the two numbers the client cared about most. AI-driven predictive analytics marketing let the team forecast campaign outcomes before launch, reallocate budgets toward top-performing campaigns, and book a 27% increase in overall marketing ROI within the first six months. Automated integration across CRM and ERP reduced manual data preparation from weeks to hours, a roughly 60% reduction in data prep time, so marketing teams can act on fresh insights almost in real time, accelerating both campaign adjustments and broader strategy execution.

Qualitative Results At A Glance

- +27% overall marketing ROI within the first six months, driven by reallocating budget toward top-performing campaigns the models could forecast in advance.
- -60% data prep time across CRM and ERP, manual preparation collapsed from weeks to hours.
- Gigabytes of daily marketing, campaign, and product data processed through a Medallion Architecture with full lineage traceability.
- Model drift no longer goes unnoticed, automated model drift detection and retraining keep predictions current without manual handoffs.
- ML-driven attribution replaced last-click reporting as the basis for budget allocation, sharpening every spend decision downstream.
- Marketing teams now act on near-real-time insights inside the same CRM and ERP they were already using, no new tool to learn.

Conclusion

Where Should Marketing And Analytics Teams Start With AI Marketing Analytics?

For this global retail analytics company, AI marketing analytics was less about adding a model and more about rebuilding the foundation that makes modeling worth doing. The Medallion data platform, optimized ETL, ML attribution, and automated model drift detection turned a fragmented, slow, opaque marketing analytics workflow into a predictive engine – one that lifts ROI 27% and gives marketing teams hours, not weeks, between question and answer.

GARRISON FLOOD CONTROL · IMPACT

What changed after the AI voice assistant went live

MANUAL DATA ENTRY PER LEAD

0

keystrokes required

Calls flow directly into HubSpot via webhook + API. Every field populated automatically by the time the call ends.

Any
fields captured
Fully configurable schema per business

~0s
CRM lag time
Webhook fires the moment the call ends

WHAT THE SALES TEAM GOT BACK

- Time back from note-taking**
Now flows into actual follow-up instead
- Full call summary in HubSpot**
Context arrives with the lead — no recordings to dig through
- Faster follow-up speed**
Lead is warm — context is already there when you call back
- No duplicate contacts**
Returning callers enrich existing records automatically
- Full audit trail in PostgreSQL**
Every call logged — replayable if the webhook ever fails

Key metric Operational gains

Retell AI · Ruby on Rails · HubSpot · PostgreSQL