



```
> execute launch_platform.sh
```

Full-Funnel AI Marketing Analytics

Replacing a \$50K/year BI stack with a deterministic, \$0/month open-source AI platform.

[dbt Semantic Layer]

[7 MCP Servers]


[Multi-Touch Attribution]

[ML Lead Scoring]


The 15-Second Miracle


Concept tag: Ask your data in plain English. Get production-grade React dashboards in seconds.

Show me the complete marketing funnel for Q1 2025: ad spend across Google and Meta, website sessions by channel, lead conversion rates, and final revenue. Calculate blended CAC and ROAS.

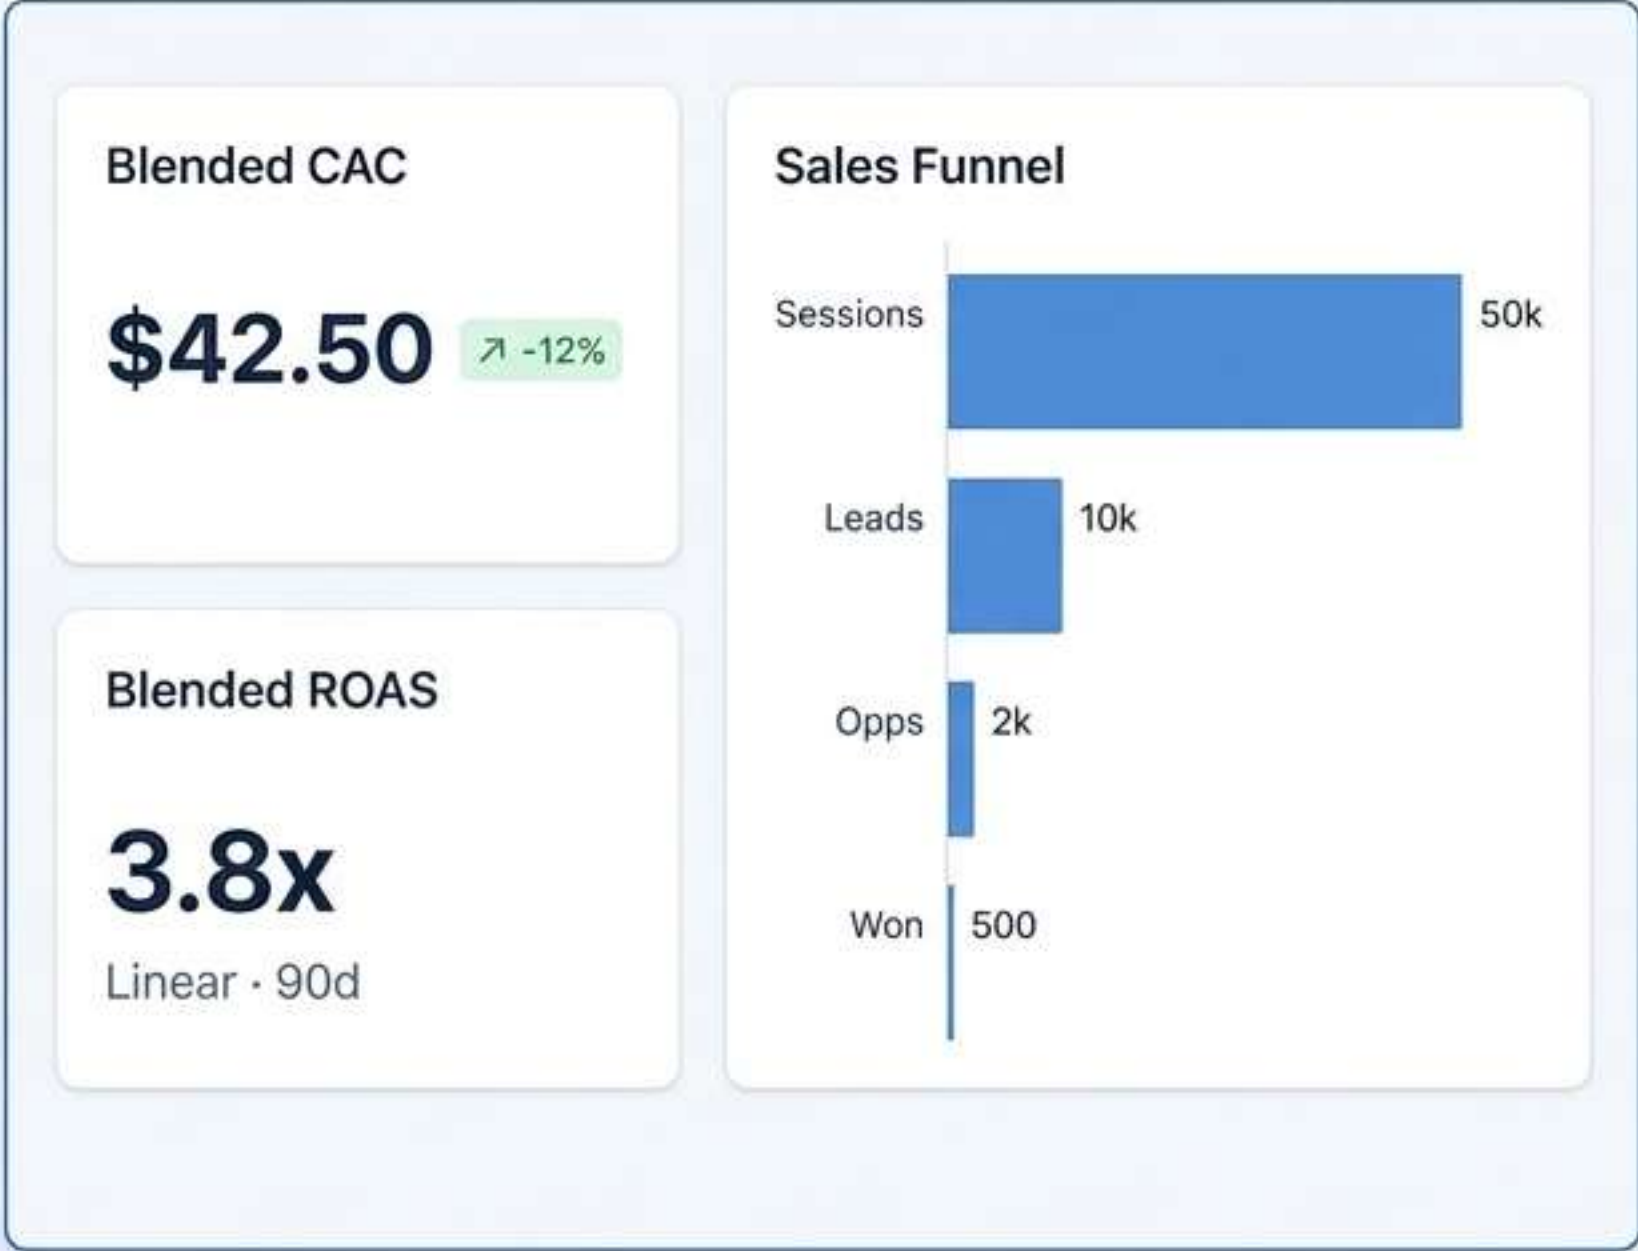


Reply...

+ 

Sonnet 4.6 

Claude is AI and can make mistakes. Please double-check responses.

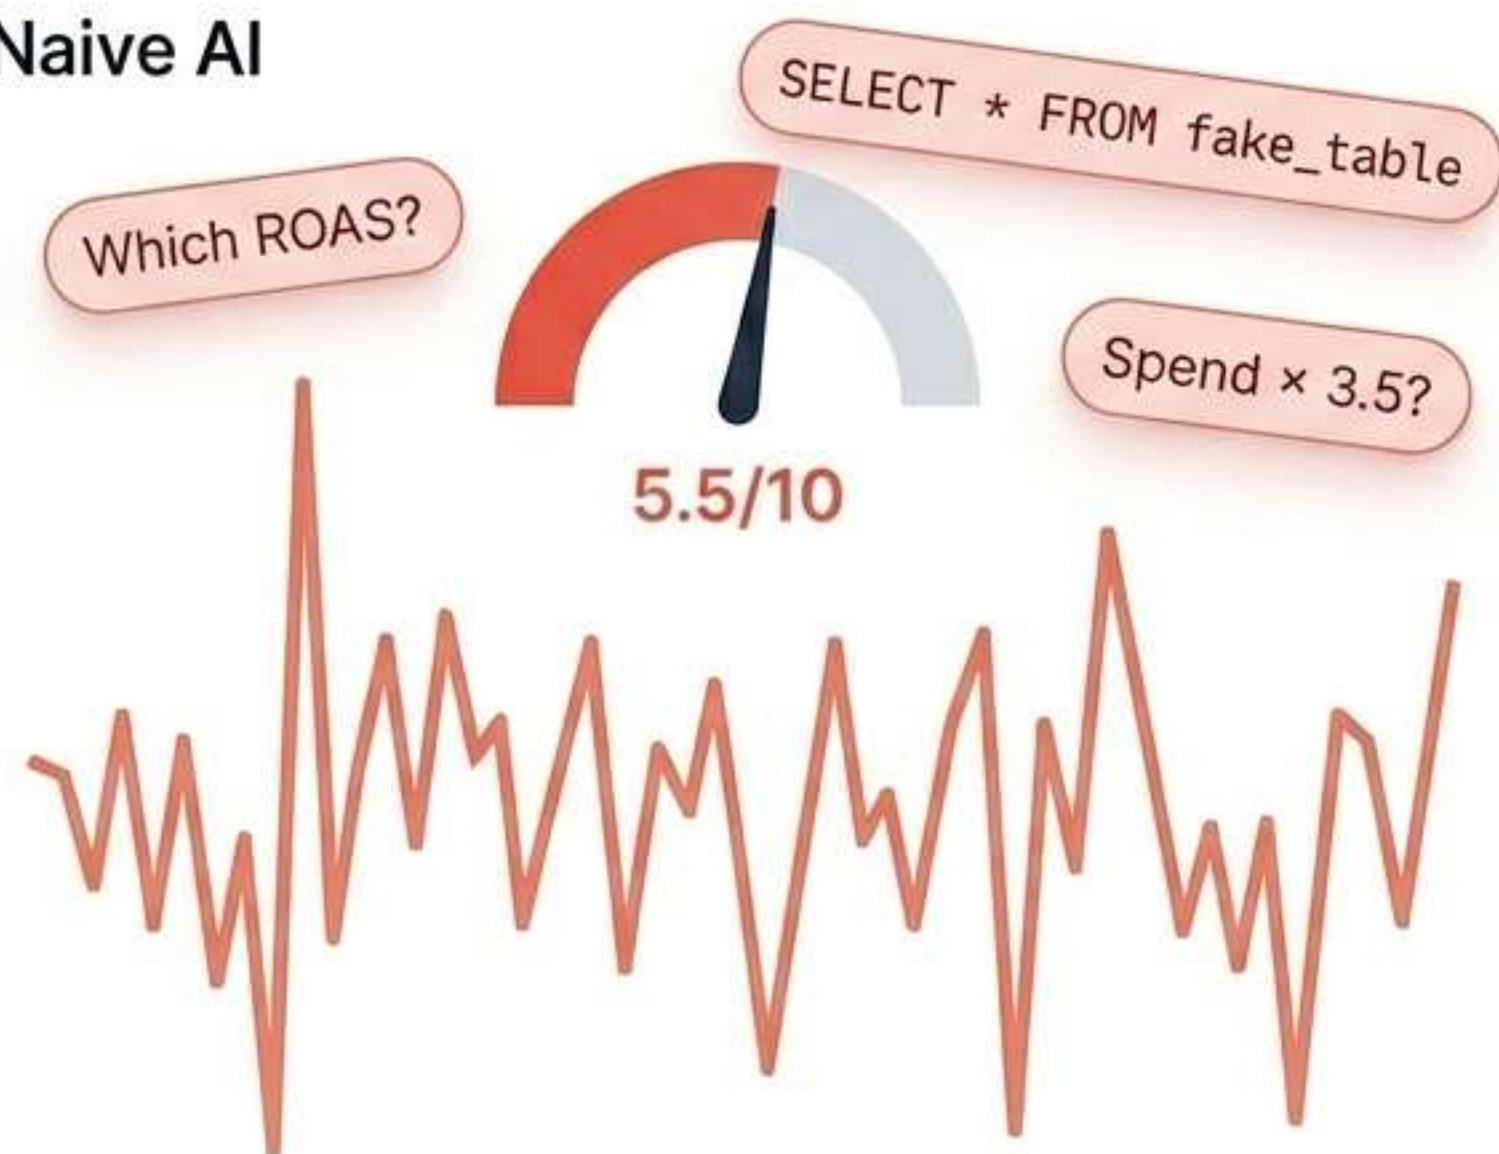


4 platforms queried. | 1 semantic layer joined. | 1 executive dashboard rendered. | Zero code written.

Why AI Analytics Fails (And How to Fix It)

Concept tag: High-end developer whitepaper Slide are two-column bento box layout.

Naive AI



Without a semantic layer, LLMs are brilliant reasoners but terrible guessers. Average confidence is just 5.5/10.

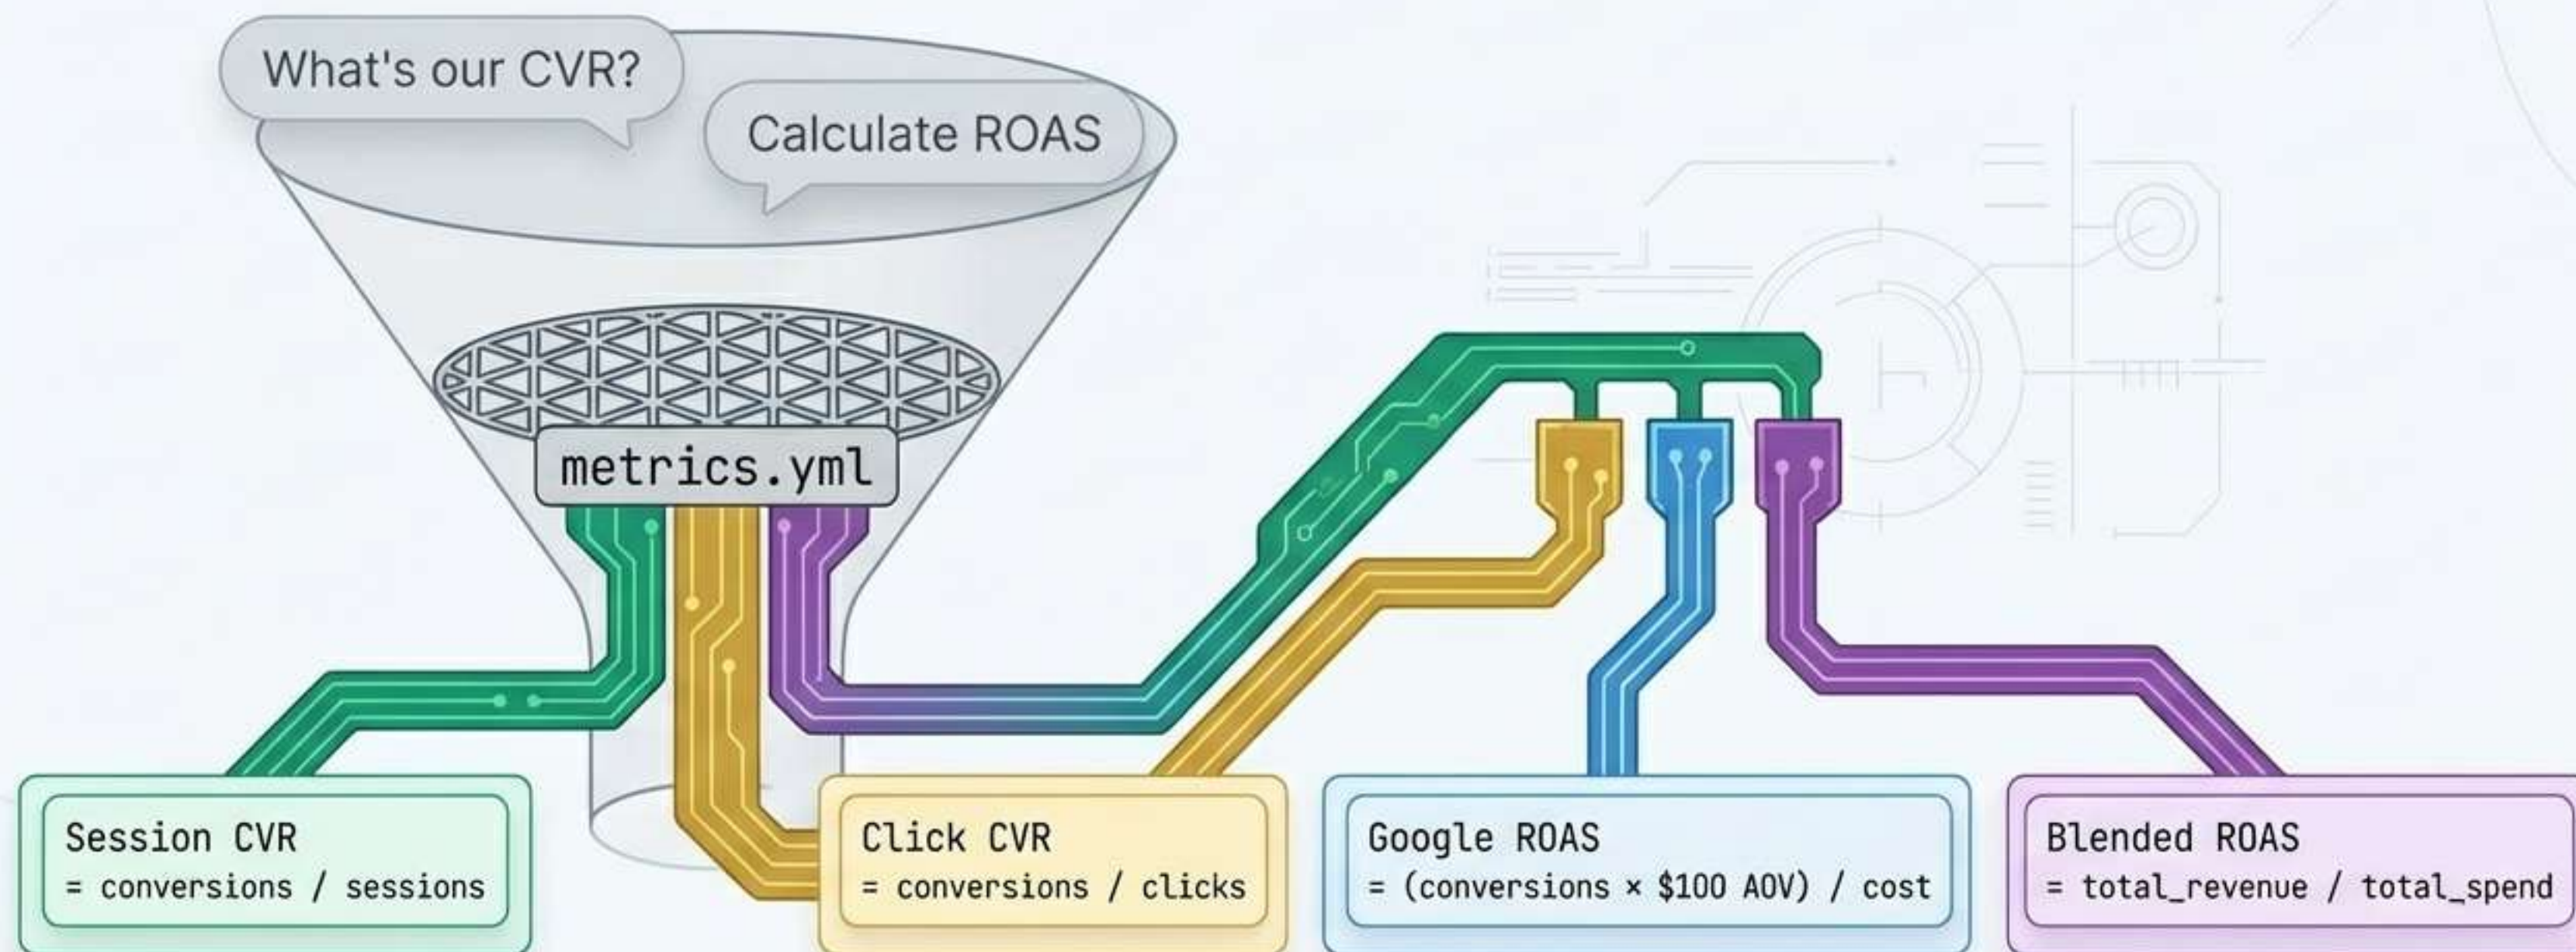
Governed AI



If metrics are strictly defined, the AI cannot invent them. 100% hallucination-free.

The Cure: A Governed Semantic Layer

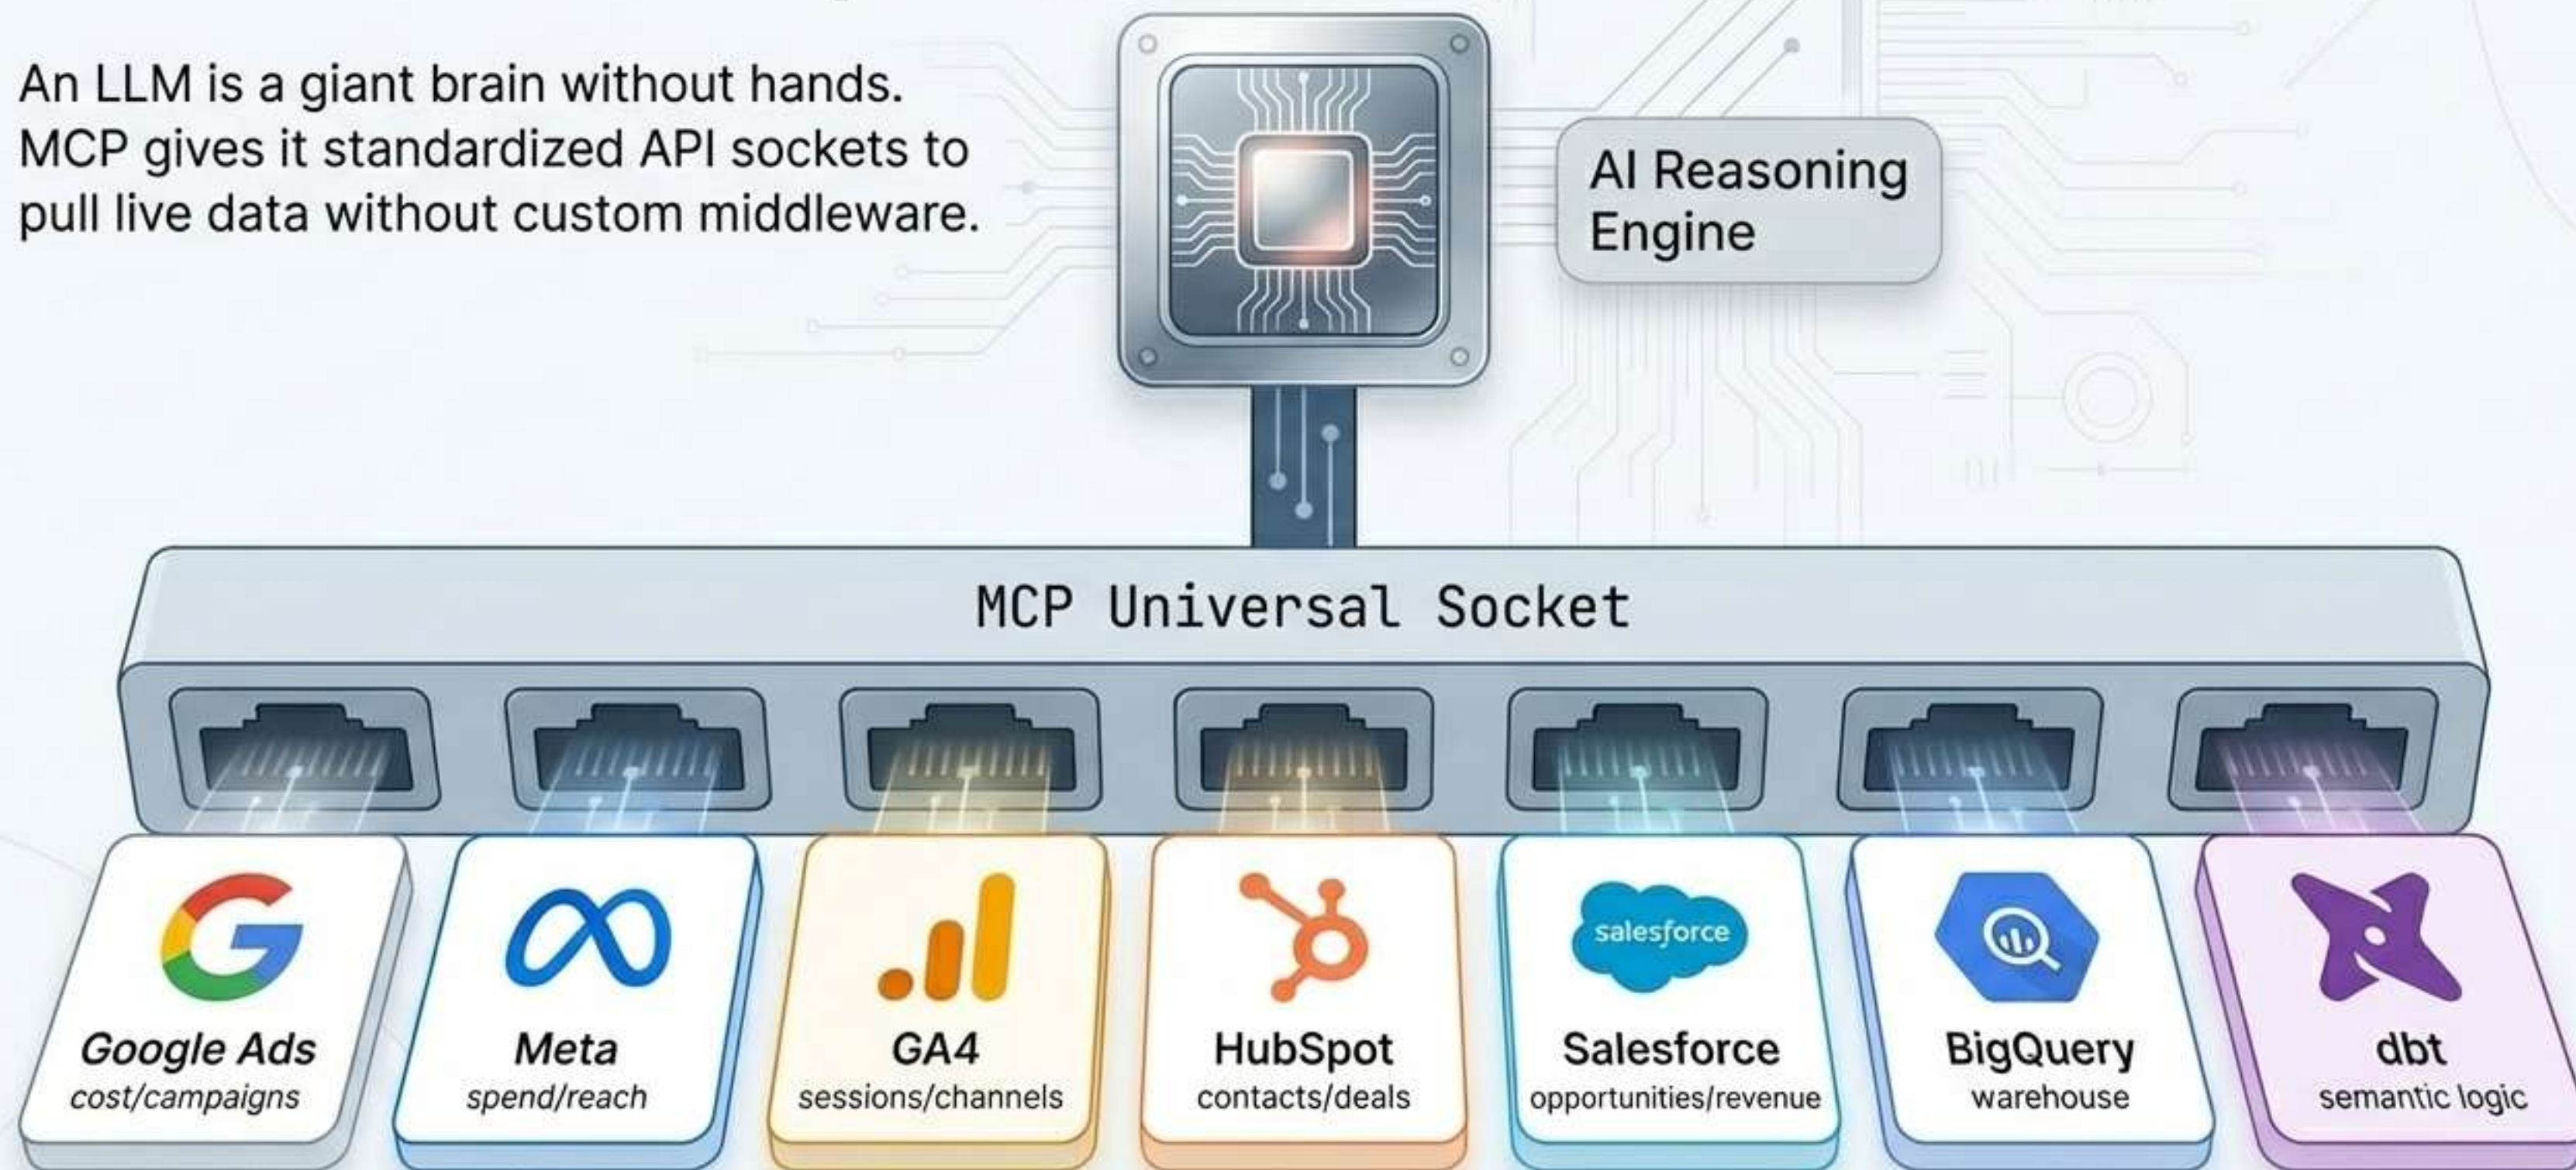
Define a metric once in YAML. Consume it everywhere flawlessly.



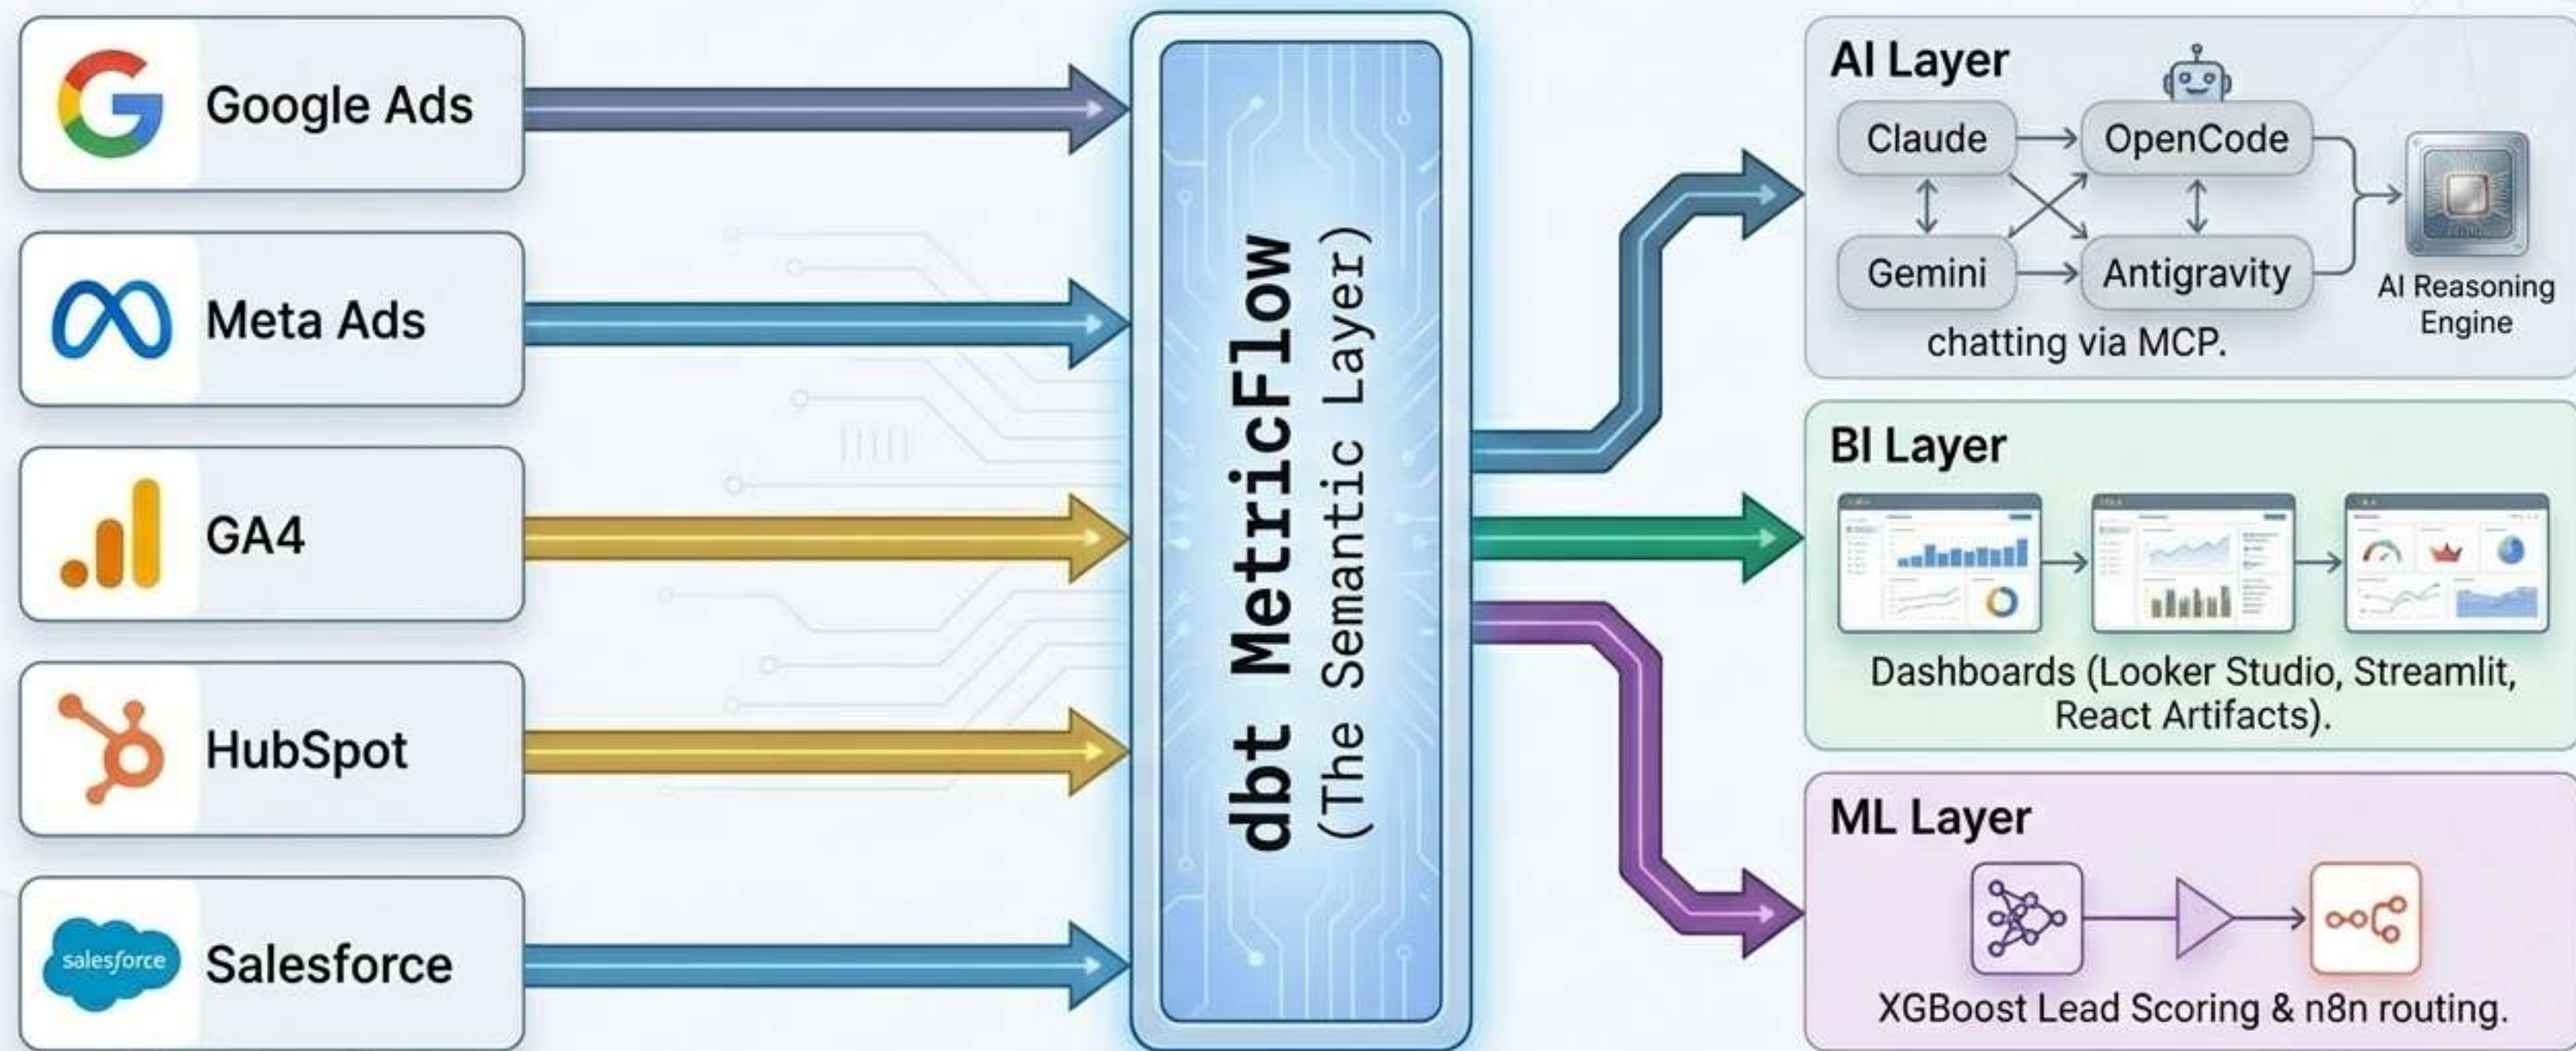
Takeaway: Never mix definitions. Never use synthetic multipliers.

Model Context Protocol (MCP): The Universal Adapter

An LLM is a giant brain without hands.
MCP gives it standardized API sockets to
pull live data without custom middleware.



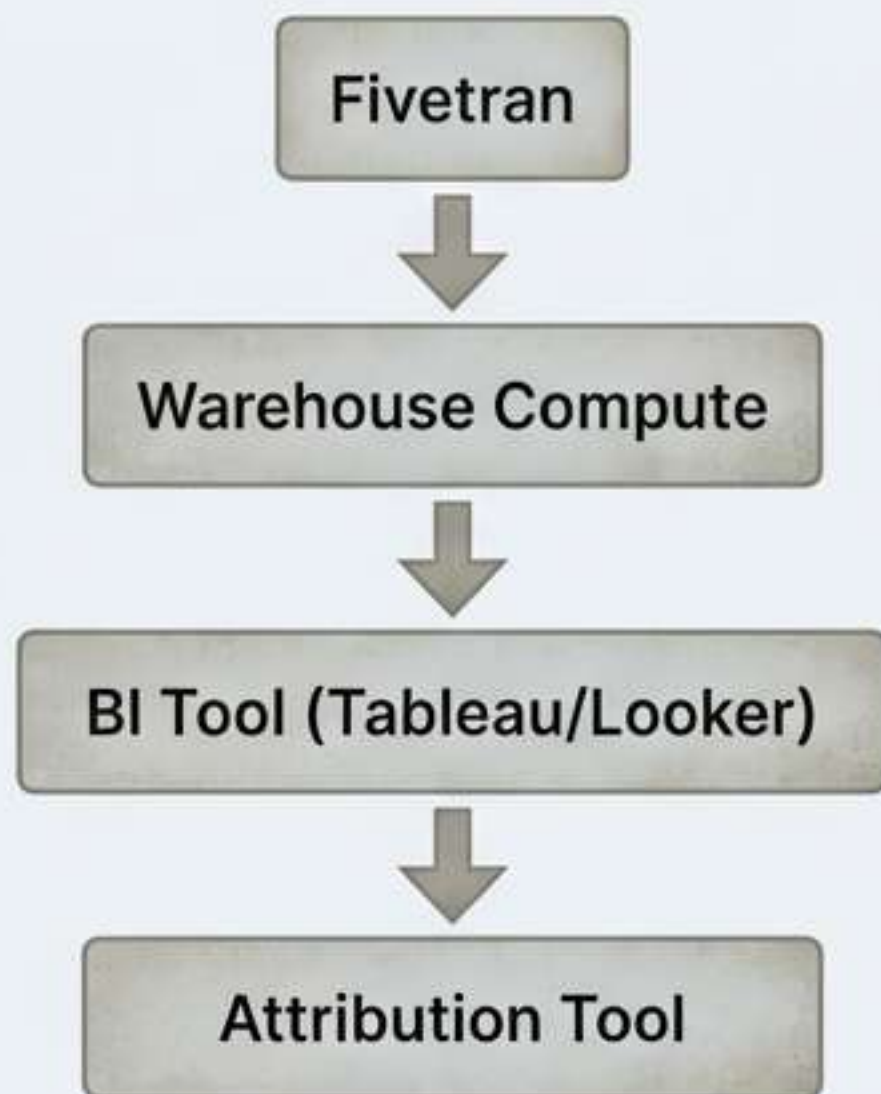
Architecture: Three Heads, One Spine



Takeaway: 2.5 million lines of code architected so every output relies on the exact same metric definitions.


The Economics of Modern Analytics

\$50k+/year



Time to insight: **Days.**

\$0/mo base cost



Time to insight: **15 Seconds.**

Zero Lock-in. Infinite Portability.

The same dbt models, the same governed metrics, and the same MCP servers run identically across all platforms. Zero code changes needed.

	Claude Desktop	OpenCode (75+ models)	Gemini CLI	Antigravity IDE
BigQuery	✓	✓	✓	✓
DuckDB	✓	✓	✓	✓
Supabase (Postgres)	✓	✓	✓	✓
Snowflake	✓	✓	✓	✓
Databricks	✓	✓	✓	✓

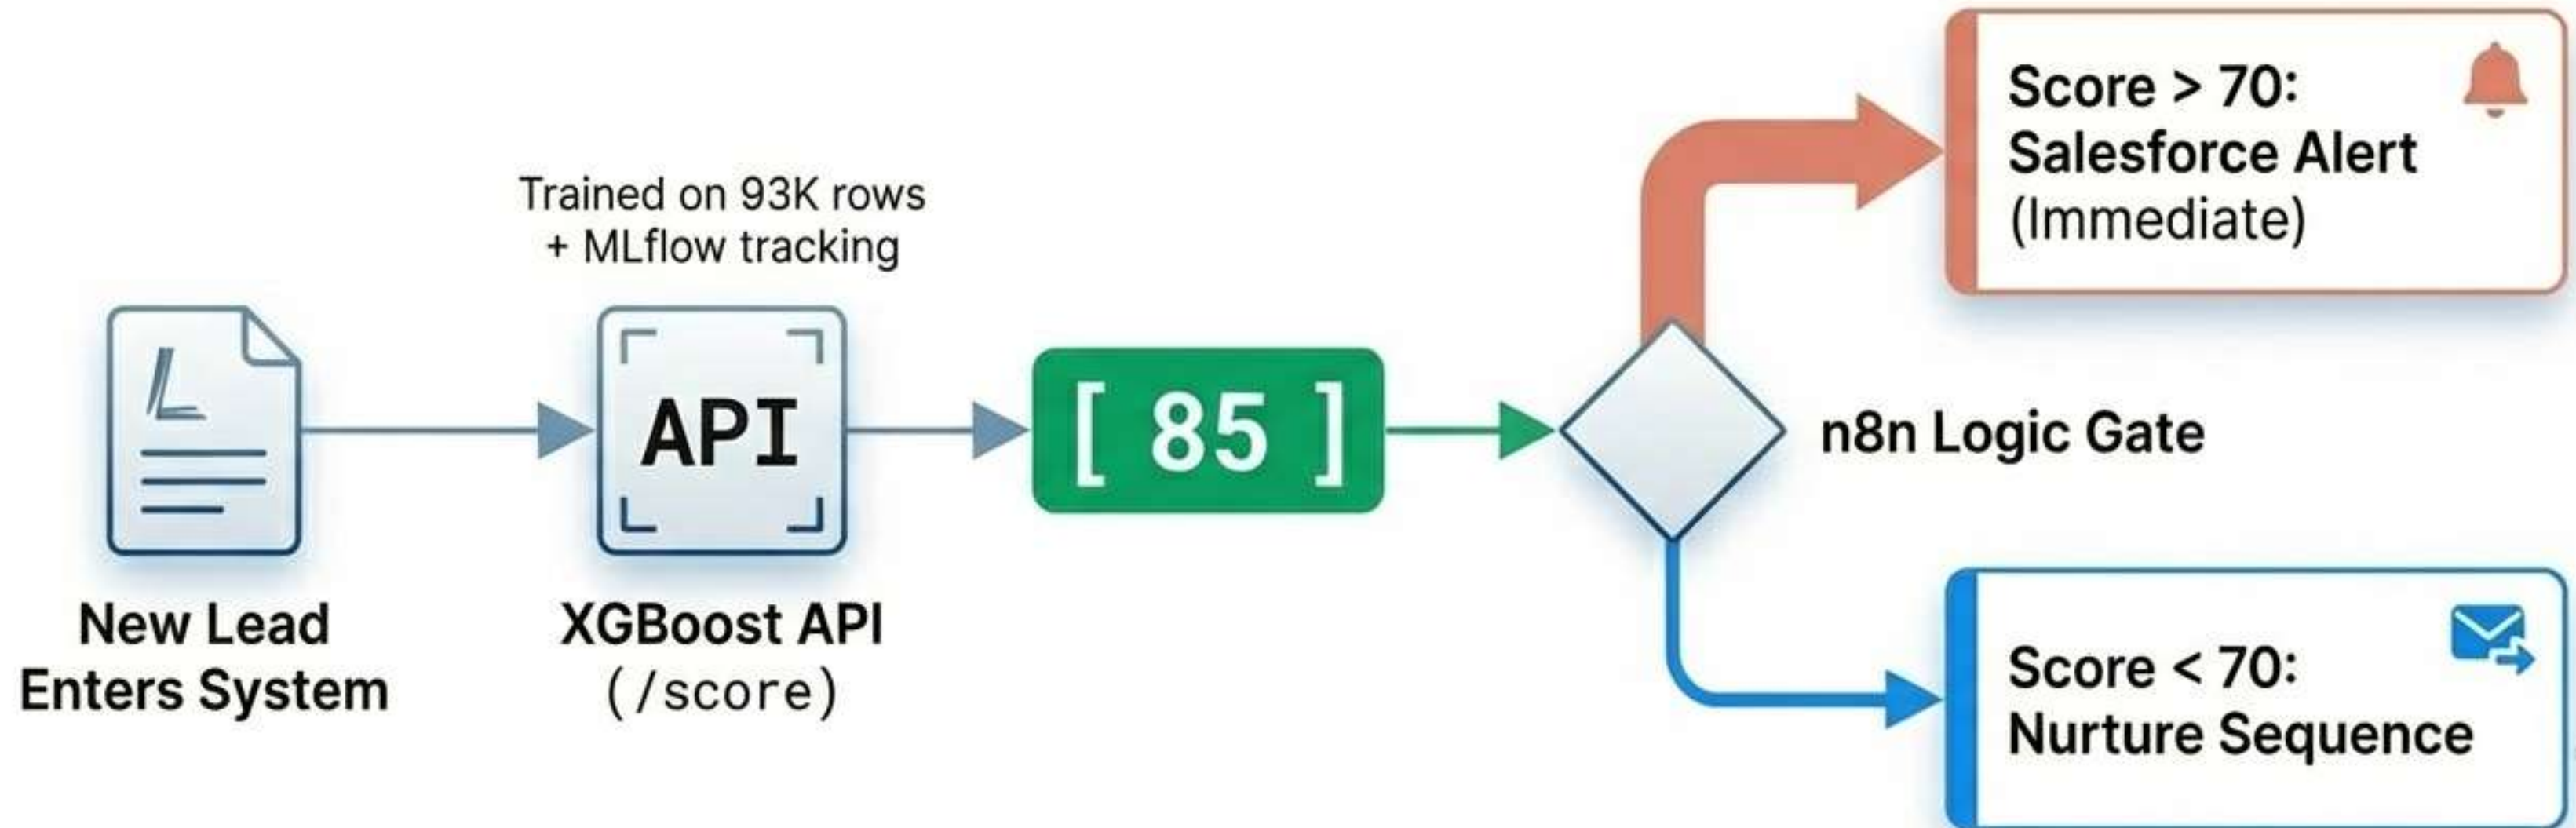
Dynamic Multi-Touch Attribution

> /attribution



Stop arguing over who gets credit. The AI calculates First-Touch, Last-Touch, Linear, and Time-Decay models simultaneously to reveal which channels drive awareness versus which close deals.

Beyond Dashboards: Automated ML Lead Routing



The all-seeing analyst also predicts your next best customer. Hot leads bypass the wait and route directly to Sales in seconds.

The Anatomy of a /marketing Command



T=0s: User types > /marketing in natural language.



T=3s: MCP intercepts. Instantly routes specialized queries to Google, Meta, GA4, and Salesforce.



T=7s: The dbt Semantic Layer translates the requests into deterministic, governed SQL.



T=12s: Live data is pulled, joined, and passed through XGBoost scoring.

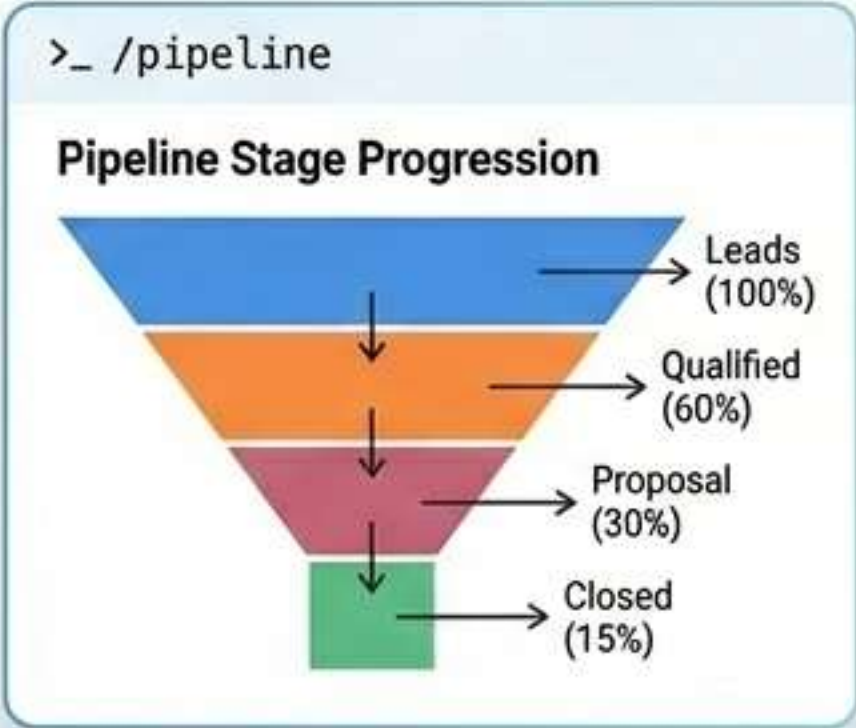




T=15s: The LLM writes and renders a complete, interactive React dashboard in the chat.



On-the-Fly React Artifacts

Every command generates a tailored, mobile-responsive dashboard using Recharts.



>_ /campaign

Google vs Meta

	CTR (Click-Through Rate)	CVR (Conversion Rate)	ROAS (Return on Ad Spend)
Google	▲ 2.5%	▲ 1.8%	▲ 3.2x
Meta	▲ 3.0%	▼ 1.5%	▼ 2.9x

The Analytics Tool Market Just Got Uncomfortable.

2.5M+
Lines of Code

2.2M+
Rows of Data

29 dbt
Models

7 MCP
Servers

5
Warehouses

Half the value prop of traditional BI tools just evaporated.
You no longer need to buy the stack; you can orchestrate it.

```
> git clone https://github.com/...
```

Ready to deploy. 100% Open Source.