Kana

A KANA WHITE PAPER

AGENTIC MARKETING

The “Just Use Claude with MCP” Fallacy

Why Claude with MCP connectors can't deliver on the promise of true agentic marketing

NOVEMBER 2026

CONTENTS

Inside this paper

· Executive Summary 3

01 The Promise of MCP 4

02 The Fragmentation Problem 5

03 The Importance of a Semantic Layer 6

04 The Cost of Reconciliation 7

05 The Need for Governed Control 8

06 What MCP Gets Right · The Limits of Skills 9

07 How Kana Closes the Gap · Conclusion 11

· Appendix: Case Study, A Premium Consumer Brand 12

· Sources 15

EXECUTIVE SUMMARY

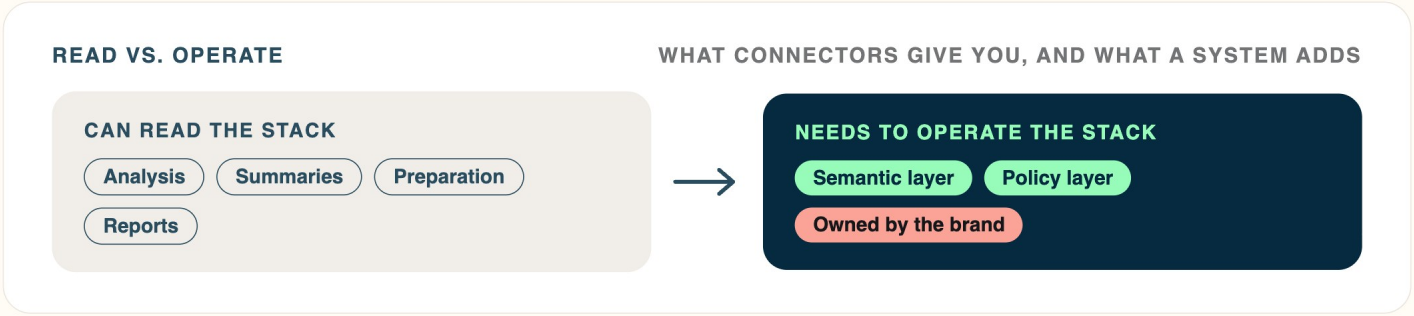
A strong model is not a marketing system

Claude is a strong model. With MCP connectors, it can read data from many ad-tech and mar-tech systems. That access does not make it a marketing system. Three structural issues stand out:

01 **Fragmentation.**
Each connector uses its own IDs, taxonomies, and audience definitions. Claude must match them on every question. Nothing keeps that match for the next question.

02 **Semantics.**
No layer stores what your brand means by “lapsed buyer,” “category,” or “incremental lift.” The same question can get a different answer on a different day.

03 **Control.**
Consent and policy rules live in each vendor's connector. No single layer enforces them. No single record shows how an agent made a decision.



The result is a capable analyst that can read the stack but cannot operate it. Brand marketers should use it for analysis and preparation. Before they let agents act, they should own a semantic and policy layer.

01 · THE PROMISE OF MCP

The Promise of MCP

Since late 2025, the major ad platforms and martech vendors have released official MCP servers:

FEB 2026

Amazon Ads

opened its MCP server to all API partners in February 2026. Agents can create, update, and report on campaigns.

APR · JUL 2026

Meta

released its ads MCP server in open beta in April 2026, and opened it to all developers in July.

APR 2026

HubSpot

made its remote MCP server generally available in April 2026, with write access.

JUL 2026

Google

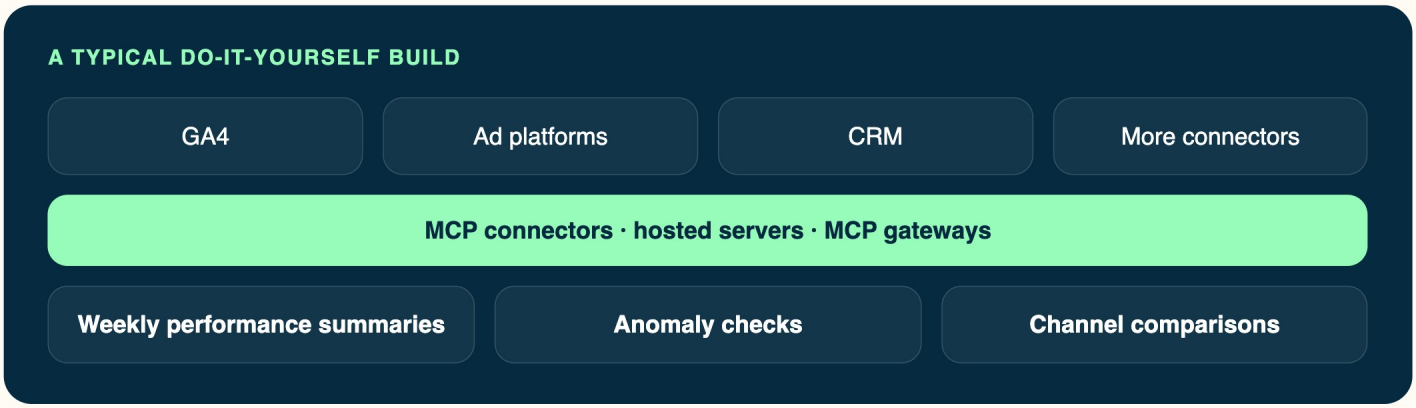
released an official Google Ads MCP server. It is read-only.

READ-ONLY

Google Analytics

released an official GA4 MCP server. It runs reports on GA4 data and is read-only.

As a result, many marketing operators now connect Claude to these servers and build their own reporting systems. A typical build pulls data from GA4, the ad platforms, and the CRM. It then asks Claude for weekly performance summaries, anomaly checks, and channel comparisons. A market of hosted servers and MCP gateways has formed to support these builds.



Each new server gives Claude access to one more system. Access to many systems is a useful start, but a marketing system needs more. This paper examines three core questions that help illustrate the difference between the Claude + MCP approach vs. what Kana delivers to customers hoping to leverage agentic marketing:

01 **Does the data agree across systems?**

02 **Does the model know what the brand means?**

03 **Who controls what an agent can do?**

02 · THE FRAGMENTATION PROBLEM

The Fragmentation Problem

An MCP connector sits in front of one system and publishes a list of actions. Claude calls the connector, and the connector decides what to return. This design is good for security, but it delivers three core issues that stand in the way of effective agentic marketing:

- **No shared model.** The CDP knows a household by one ID. The retail media network uses a hashed email. The syndicated panel uses a panelist ID and projects to the market. The ad platform reports by campaign. When a question spans these systems, Claude must match the records while it answers. The next question starts again from zero.
- **Every question is an integration project.** A question such as “Which audiences drove category growth last quarter?” can call five connectors. Claude pulls the results into the conversation and resolves the conflicts. The marketer cannot see those decisions. When the CDP and the retail media network disagree on audience size, which number did Claude use?
- **The vendor sets the limit.** Each connector exposes only the actions its owner allows. Many marketing connectors are read-only or expose reports only. Claude can describe an audience. Often it cannot build that audience in the platform, schedule it, or send it to activation.

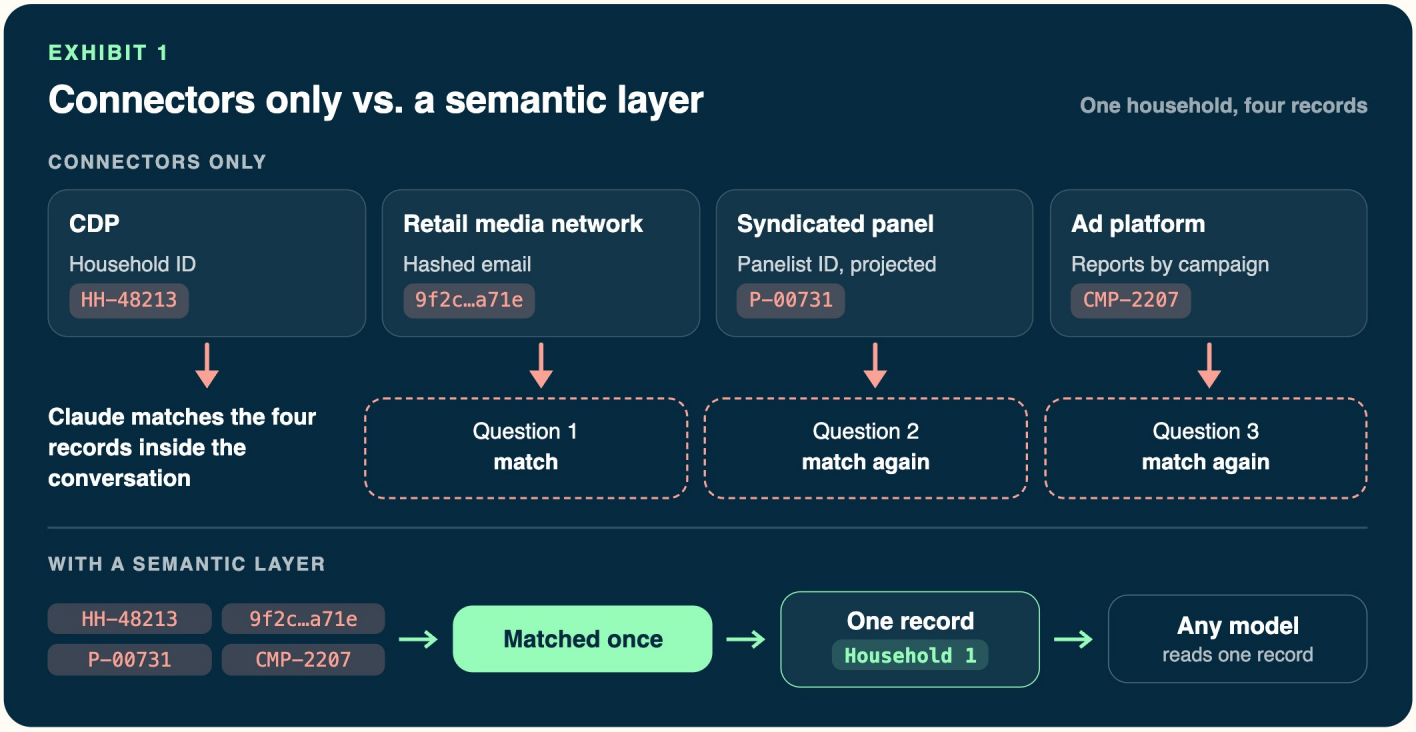


Exhibit 1. *On top, one household appears four ways, and Claude matches them on every question. Below, the match happens once, and any model reads one record.*

03 · THE IMPORTANCE OF A SEMANTIC LAYER

The Importance of a Semantic Layer

Matched IDs are only the first step. Marketing data also needs shared meaning. Each brand has its own definitions: what counts as a lapsed buyer, which products are in the category, which time window measures lift. These definitions live in analysts' heads, in slide footnotes, and in SQL that few people read.

With connectors only, Claude infers a definition each time it answers. It uses the data it receives and the words in the prompt. A small change in either can change the definition. The answer stays fluent, so the change is hard to see.

The problem grows with each new user:

DIVERGENCE

Two users can ask the same question through the same connectors. Claude can make different reconciliation choices for each user, so they get different answers. The CMO then sees two numbers in two decks.

NO SHARED LEARNING

When one analyst corrects a definition or a join, the fix stays in that analyst's conversation. Other users do not get it. Projects and skills can share some fixes, but they bring the costs described in “The Limits of Skills.”

Exhibit 2 shows the effect. Three people ask the same question through the same connectors, and they get three answers. Each answer is plausible. The CMO cannot tell which one to take to the board.

EXHIBIT 2

“How many lapsed buyers does Brand X have?”

Illustrative values

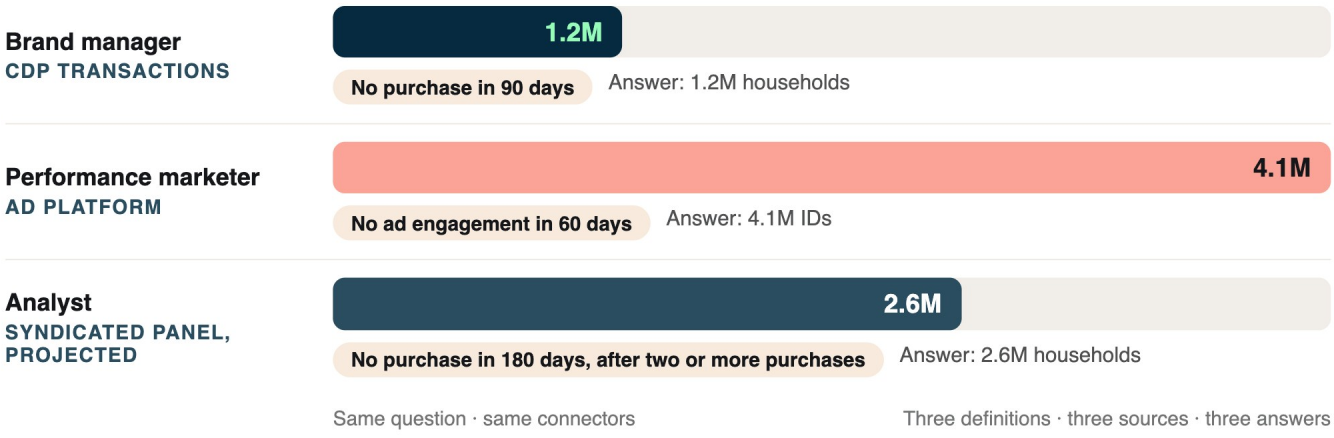


Exhibit 2. *“How many lapsed buyers does Brand X have?” (illustrative values). Each row shows who asked, the definition Claude used, the source Claude trusted, and the answer.*

A semantic layer fixes this upstream. It stores each definition once, with an owner and a version. Every agent and every analyst uses the same definition. The answer changes only when the data changes.

04 · THE COST OF RECONCILIATION

The Cost of Reconciliation

Each question that Claude answers through connectors has a token cost, and this architecture makes that cost higher than it needs to be. The cost comes from four places:

01

Raw results in context.

Each connector call puts raw rows into the conversation. In this pattern, input tokens are usually the largest part of the bill.

02

Reconciliation.

Claude uses more tokens to match IDs and resolve conflicts between systems.

03

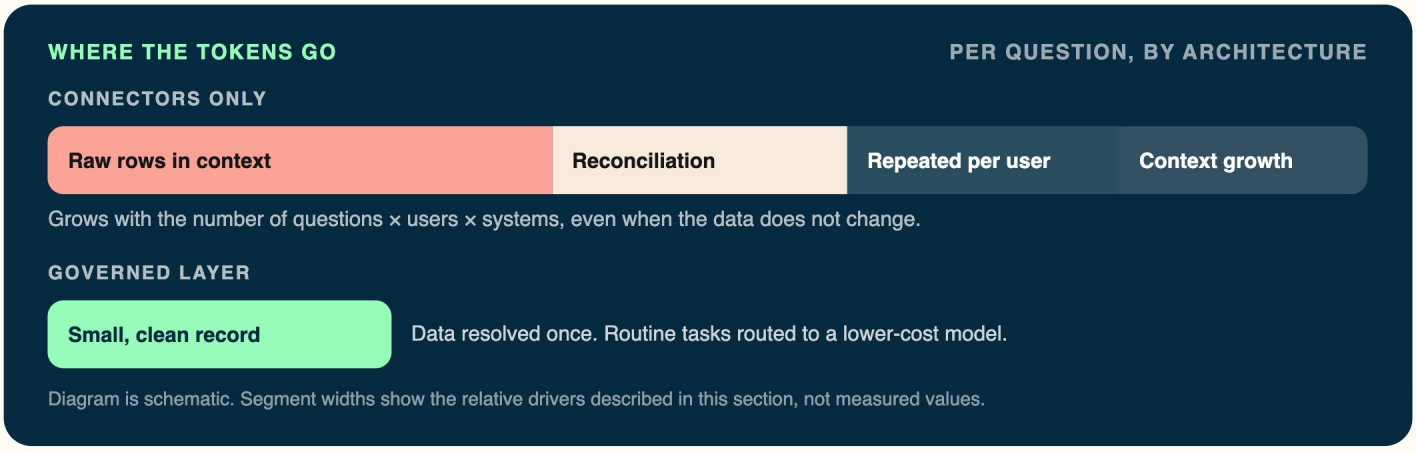
Repetition.

Nothing persists between questions. The same reconciliation runs again on the next question, and again for each user.

04

Context growth.

In a long conversation, the model reads earlier results again on each turn. Loaded skills add more tokens.



The cost grows with the number of questions, users, and systems. A team that asks more questions gets a higher bill, even when the data does not change.

A governed layer reduces this cost. When the data is resolved once, the model reads a small, clean record. Kana also picks the model for each task, so routine tasks use a lower-cost model, ensuring you don't keep paying to do the same analysis or make the same decisions over and over again.

05 · THE NEED FOR GOVERNED CONTROL

The Need for Governed Control

An insight has value only when it becomes an action. In marketing, an action is an audience built, a consent check passed, a campaign scheduled, and a result measured. Each step lives in a different system, and each system has its own rules.

Consent is fragmented. Each connector enforces the consent signals its owner holds. No connector knows the full set of consent and purpose rules for a person. In one conversation, an agent can join data that no single system would let it join.

Policy has no home. Brand rules, regulatory limits, and partner contract terms sit outside every connector. Examples: no health inference in pharma audiences, and no retailer data outside that retailer's network. Claude follows these rules only when the prompt tells it to.

No audit trail. When an agent recommends or builds an audience, the brand must know which data, definition, and rule it used. A conversation log is a weak record for a regulator or a partner audit.

Exhibit 3 shows where the chain breaks:

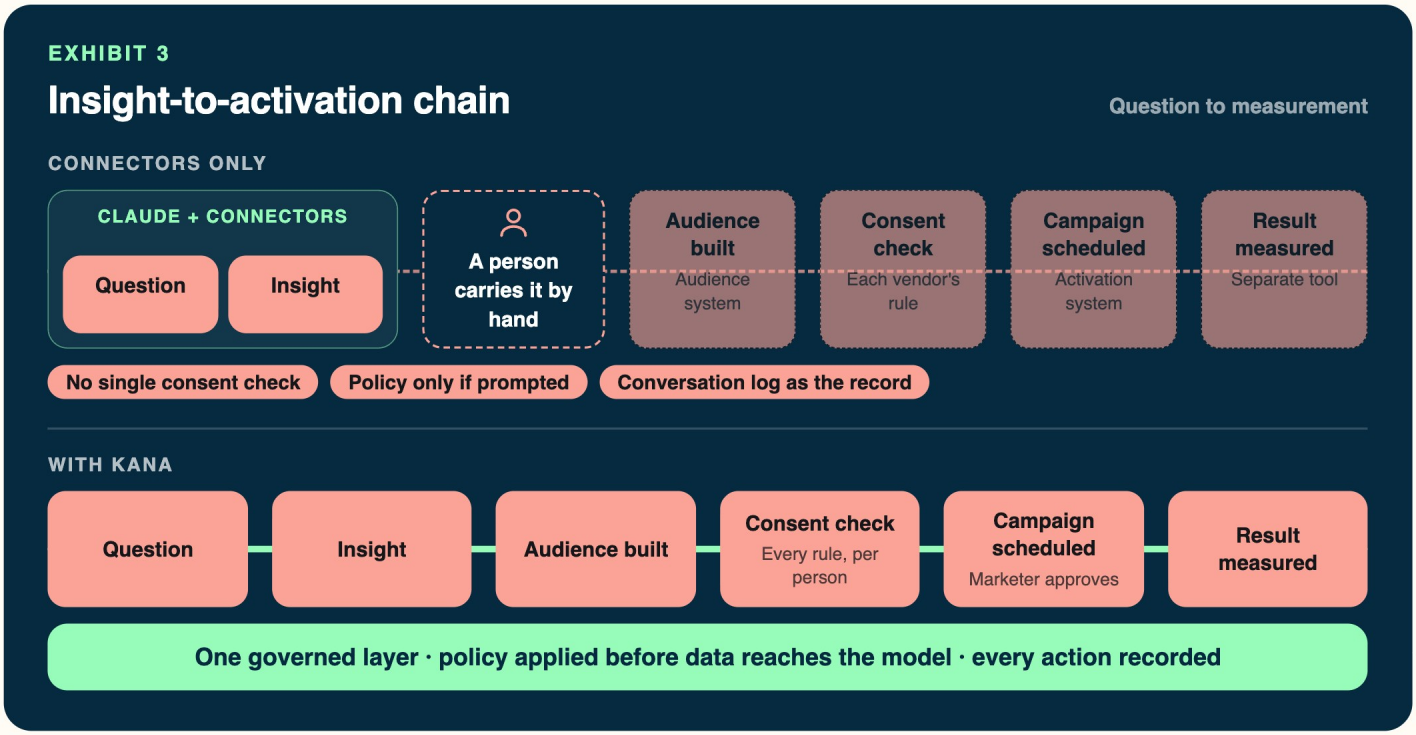


Exhibit 3. *With connectors only, a person must carry the insight into the audience and activation systems by hand. With Kana, one governed layer carries it from question to measurement.*

06 · WHAT MCP GETS RIGHT

What MCP Gets Right

In the spirit of fairness and accuracy, it's important to note what this pattern gets right:

The model is strong.
Claude reads, summarizes, and explains data well. For campaign reviews and briefs, that quality shows.

MCP is open.
Any vendor or brand can build a connector. Connectors move between models.

No migration.
Connectors work with the systems a brand already owns.

Read-only is a safe start.
An agent that cannot spend media budget cannot waste it.

The issue is the architecture, and the distinction is that this doesn't deliver on the true promise and potential of agentic marketing.

THE LIMITS OF SKILLS

The Limits of Skills

On the surface, skills may seem to be an effective way to plug the gap. A skill is a set of instructions and scripts that Claude loads when a task needs it. A brand can write skills that state its definitions, its matching rules, and its policy rules.

Skills help. A written definition of “lapsed buyer” reduces drift. A skill that calls a stored ID crosswalk reduces the reconciliation work.

That said, a skills library is also software, and software has costs:

- **Authoring.** Someone must write each definition, matching rule, and policy rule for each brand and each source.
- **Maintenance.** Connectors, taxonomies, product lists, and consent rules change. Each change can break a skill with no warning, so someone must update and re-test the skills continuously.
- **Evals.** Each skill needs test sets, expected answers, and a regression run when the model or a connector changes.
- **Governance.** Each skill needs an owner, a version, a review, and a rule about who can change it.

This work requires technical resources: data engineers, analytics engineers, and people who can build and run evals. MCP connectors supply none of them. Most marketing teams do not have these people, and the data team has other priorities. The cost of a skills library continues for as long as the brand uses it.

06 · THE LIMITS OF SKILLS, CONTINUED

Even a well-run library leaves gaps. A skill tells the model what to do, but it cannot stop the model from doing something else. It cannot add write actions to a connector, hold a resolved identity, or produce an audit record. In effect, a skills library is a semantic layer in markdown, with no enforcement.

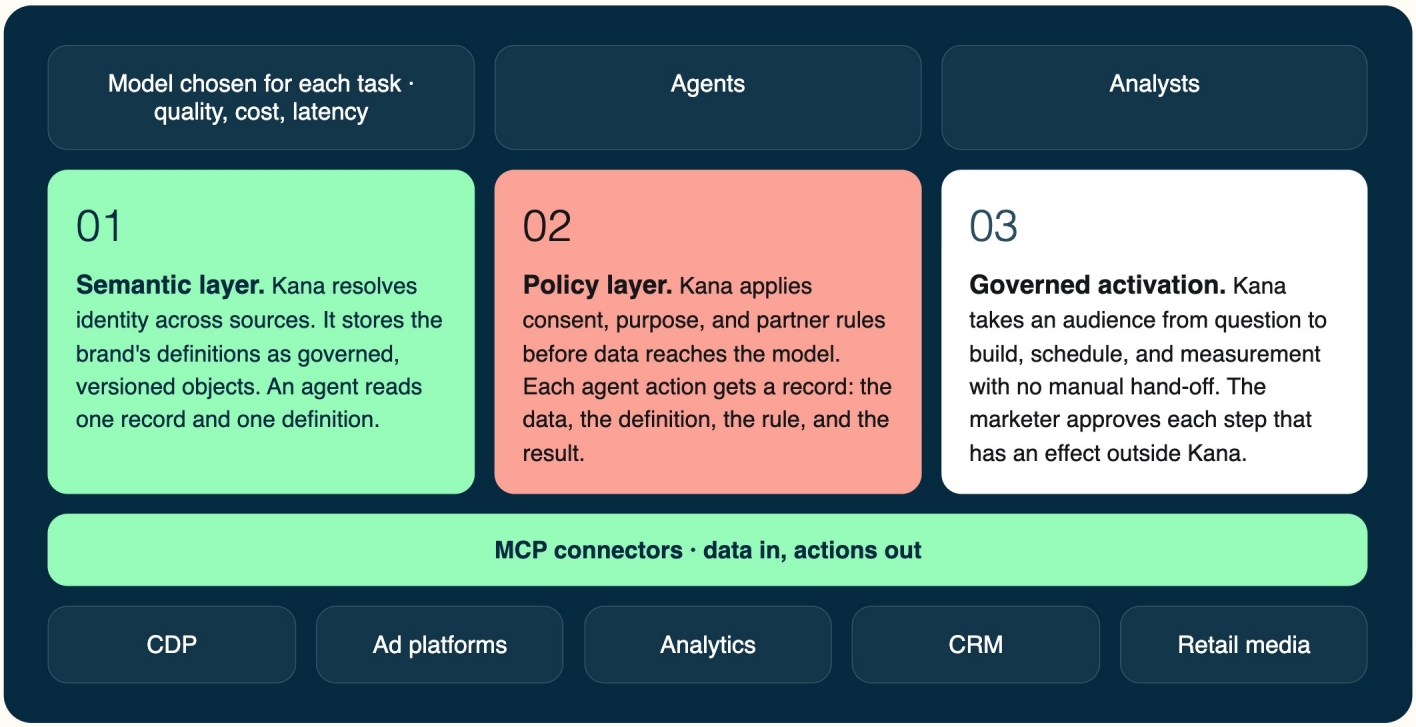
WHAT A SKILL CAN AND CANNOT DO

What skills can do	What skills can't do
✓ Tell the model what to do	✕ Stop the model from doing something else
✓ Reduce drift and reconciliation work	✕ Add write actions to a connector
	✕ Hold a resolved identity
	✕ Produce an audit record

07 · HOW KANA CLOSES THE GAP

How Kana Closes the Gap

Kana is the layer between the marketing stack and the model. It does three things that connectors cannot do.



Kana also uses MCP. Connectors bring data in and send actions out. Kana supplies the shared meaning and control between them. Kana picks the best model for each task, based on quality, cost, and latency. The brand does not have to track which model is best.

CONCLUSION

Conclusion

Claude with MCP connectors gives brand marketers a capable analyst with access to the whole stack. That analyst reads through connectors it does not control. It guesses at definitions. It works under consent rules that nobody enforces in one place. This is useful for analysis. Agents that act for a brand need more.

> “
>
> The brand that wins with agents will own its definitions, its consent rules, and its record of decisions. With Kana, the brand also gets the best model for each task, with no model decisions to make.

APPENDIX · CASE STUDY

Case in Point: A Premium Consumer Brand

This case study comes from a Kana customer engagement. Some details are changed to protect the customer's confidentiality. It is written for marketing operations and data teams.

Context

The brand profiled in this case study makes premium consumer products. It sells through three channels: DTC e-commerce on Shopify, its own retail stores, and independent dealers, which include large key accounts. The marketing team wants four reporting areas, each for one user group:

Executive marketing impact for the CMO and CEO

Campaign ROI for brand and performance marketers

Portfolio performance for commercial and product leads

Marketing operations for program managers

All four depend on one shared data and taxonomy foundation. Two gaps exist before any AI work starts. In-store sales need a CRM build before they can link to customers. The product taxonomy is being rebuilt.

System Map

EXHIBIT 4

The brand's seven data sources

No two share a customer ID

SOURCE	WHAT IT HOLDS	KEY ID	GRAIN
Shopify	E-commerce orders, customers, products	`Shopify customer ID` `order ID`	Order line
GA4	Site behavior and traffic source	`GA4 client ID (cookie)`	Session and event
Meta Ads	Paid social spend and delivery	`Campaign, ad set, and ad IDs`	Ad by day
Google Ads	Paid search, Performance Max, and video	`Campaign ID` `click ID`	Campaign by day
ShopSense / Locally	Dealer sell-through and stock	`Dealer location` `SKU`	Store by SKU by period
Emarsys	Email and SMS contacts and engagement	`Emarsys contact ID` `email`	Contact by send
Sprout Social	Organic posts, engagement, sentiment	`Profile and post IDs`	Post

Shopify ↔ GA4: order ID, only when tagging works **Dealer data: no customer at all**

Exhibit 4. *The brand's seven data sources. No two sources share a customer ID. Shopify and GA4 join through the order ID, but only when tagging works. Dealer data has no customer at all.*

APPENDIX · CASE STUDY, CONTINUED

Four Questions

EXHIBIT 5

Four questions from the brand's reporting plan

Question · break · resolve

QUESTION	WHERE CONNECTORS BREAK	WHAT KANA RESOLVES ONCE
What did marketing contribute to revenue across DTC, stores, and dealers? Shopify · GA4 · Meta Ads · Google Ads · ShopSense / Locally	Store sales do not link to customers yet. Dealer data has no customer. Claude picks an attribution rule each time.	One revenue definition for each channel, and one versioned attribution method
What is our iROAS on Meta compared with Google? Meta Ads · Google Ads · GA4 · Shopify	Each platform reports its own attributed revenue. Campaign codes differ. “Incremental” has no shared definition.	One campaign taxonomy and one iROAS definition with a stated method
Which persona segment drives premium mix? Shopify · Emarsys · GA4 · Sprout Social	No source holds the persona. “Premium” depends on a product taxonomy that is still being rebuilt.	Persona and premium tier stored as governed definitions on one customer and one product record
Did the video campaign lift dealer sell-through? Meta Ads · Google Ads · ShopSense / Locally · Sprout Social	Ad data is by ad and day. Dealer data is by store and SKU. Claude must invent a match by place and time.	One crosswalk from ad delivery to dealer locations, and one lift method

Exhibit 5. *Four questions from the brand's reporting plan, the systems each one needs, where connectors break, and what Kana resolves once.*

> “Each question requires access to four or more systems to answer. With connectors only, Claude rebuilds each join on every request. The answer depends on choices that the marketer cannot see.”

APPENDIX · CASE STUDY, CONTINUED

Activation and Consent

The four questions above ask for insight. The next request asks for action:

> “
>
> **Retarget site visitors who viewed premium models but did not buy, through email and paid social.”**

This one request touches five systems, and each system holds a different consent rule for the same person.

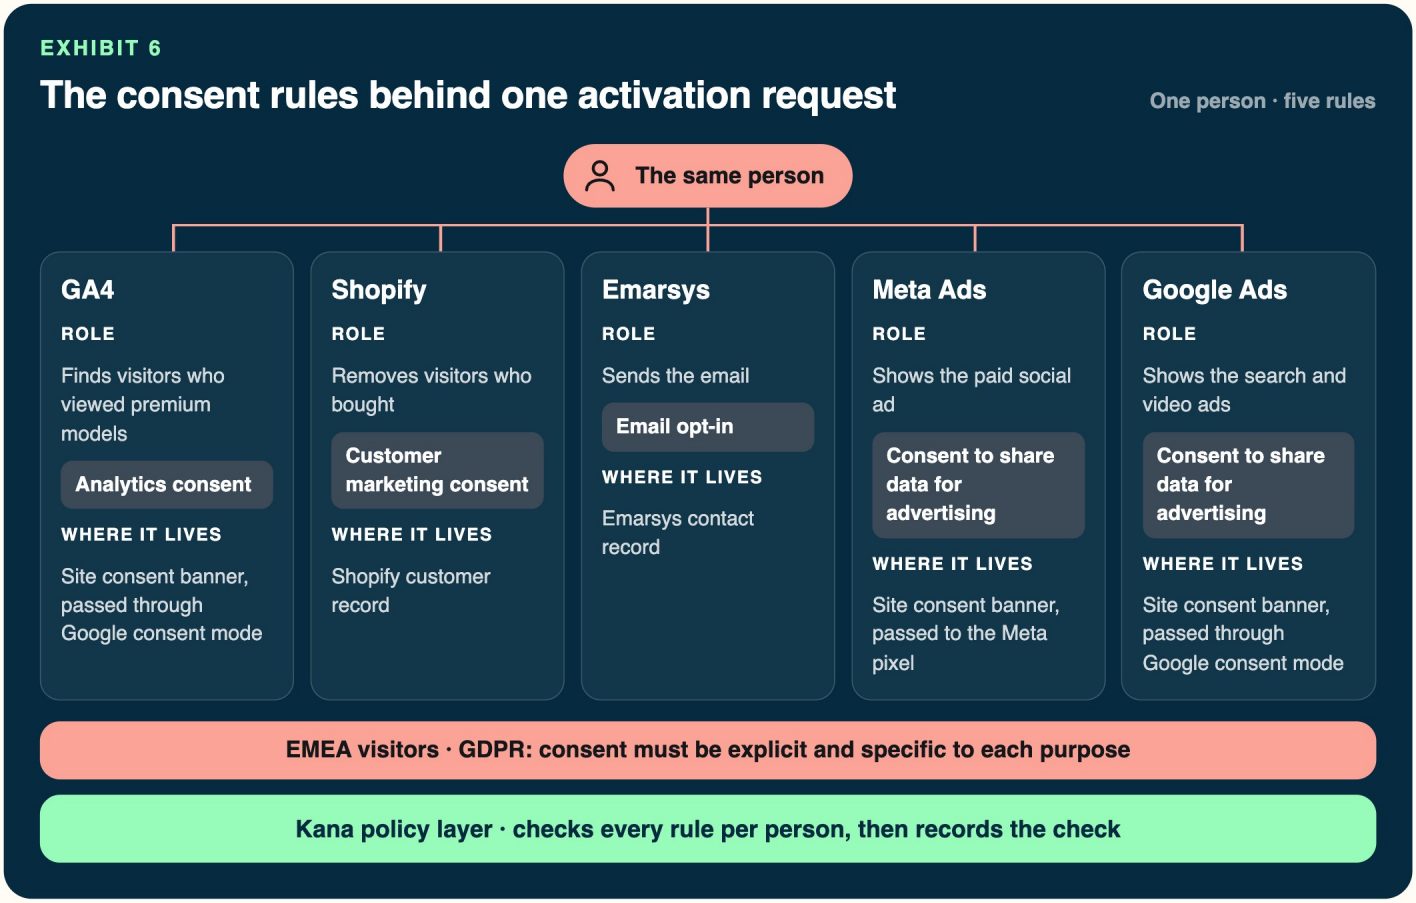


Exhibit 6. *The consent rules behind one activation request*

For visitors in EMEA, GDPR adds one more condition: consent must be explicit and specific to each purpose.

No single connector sees all five rules. With connectors only, Claude must join five consent states for each person inside the conversation. If one state is missing or old, the agent can email or target a person who did not agree. Kana's policy layer checks every rule for each person before it builds the audience, and it records the check.

SOURCES

1. Amazon Ads. Amazon Ads MCP server open beta. advertising.amazon.com/library/news/amazon-ads-mcp-server-open-beta
2. Meta ads MCP server, open beta. adsuploader.com/blog/meta-ads-mcp
3. Meta for Developers. Meta Ads MCP server, July 16, 2026. developers.facebook.com/blog/post/2026/07/16/meta-ads-mcp-server/
4. Official Google Ads MCP server. admakeai.com/connect/google-ads-to-claude
5. HubSpot remote MCP server. claap.io/blog/hubspot-mcp
6. Official GA4 MCP server. crox.io/servers/google-analytics

FURTHER READING FROM KANA

BLOG

Individual Agents vs. Enterprise Agents

Read the article →

BLOG

The Build Is Not Your Job

Read the article →

GUIDE

Evaluation Guide: Is Your AI Enterprise Grade?

Get the guide →

ABOUT KANA

Founded in 2025 by Tom Chavez and Vivek Vaidya (the martech pioneers behind Krux, Rapt, and Habu), Kana delivers agentic marketing applications custom-built to your goals and architecture with enterprise-grade scale and security. Backed by Mayfield. Based in San Francisco. Learn more at www.kana.ai.

Ready to get started?

Talk with our team about owning the semantic and policy layer before your agents act.

Talk to our team →