

— A KANA WHITE PAPER

CONSUMER PACKAGED GOODS

Getting Ahead of the Category

Why your category data keeps telling you what already happened, and how agentic marketing will change how CPG marketers go to market in the near future.

CONTENTS

Inside this paper

—	Executive summary	3
01	CPG’s timing problem is an architecture problem	4
02	What an agent is, and what it is not	6
03	Addressing the readiness gap	8
04	What enterprise grade has to mean in CPG	10
05	The build is not a brand marketer’s job	13
06	What agentic marketing means when applied to CPG	15
07	What a build-with agentic marketing engagement requires	17
08	Conclusion	18
—	Endnotes	19

EXECUTIVE SUMMARY

You get the data, but weeks after it can be put to use

Let's begin with an unfortunate historical fact about consumer packaged goods. In the past, you've been compelled to buy information about that category back from somebody else, and sometimes weeks after it stopped being true. The retailer owns the transaction. The syndicated provider owns the read. You own the brand, the P&L, and the risk, and you learn how the quarter went from a report describing what's already ancient history by that point.

That was survivable when categories moved slowly and shelf was the primary battleground. It isn't now. Retail media is fragmenting faster than anyone's measurement can follow, and AI assistants have started deciding which product a shopper hears about based on which one gives the best answer, not which one has the biggest brand. A century of brand equity is being repriced by systems that were never taught to care about it.

This paper outlines why Agentic AI applied to your marketing data can help finally close the gap between when something becomes true in your category and when you can act on it. It's not a faster dashboard and it's not a better chatbot. It's a different architecture, one in which software pursues a goal, plans the steps, acts across your systems, evaluates what happened, and adapts, all inside the policies you set. The act of agentic marketing applies that principle to your marketing and go-to-market strategy.

This white paper is structured to cover four topics:


- 01 What an AI agent actually is and how it differs from the automation and generative tools already in your stack
- 02 What the evidence says about how ready CPG organizations really are to adopt agentic AI
- 03 What enterprise grade has to mean when your data lives across a warehouse, a dozen retail media networks, and a syndicated subscription
- 04 How agentic marketing applies to the specific decisions your category and brand teams make every week

We wrote this paper for CPG marketing leaders who've been sold transformation before and want to understand what's structurally different this time around. Let's dive in.

01 • THE TIMING PROBLEM

CPG's timing problem is an architecture problem

The instinct when marketing feels slow is to buy something that promises to speed it up. That instinct is largely responsible for creating the current problem facing marketers in CPG companies. In 2011 there were roughly 150 martech solutions. Today there are more than 15,000.¹ Most CPG teams have accumulated at least a CRM, a marketing automation platform, a warehouse, a syndicated data subscription, an agency relationship, and a growing pile of retail media logins that seem to reconcile with nothing.


THE STACK GREW. THE CONNECTIVE TISSUE DIDN'T.

Source: Kana, *Understanding AI Agents for Marketing*.¹

What was never built is the connective tissue that makes them work as a system, so your team became that connective tissue. As a direct result, experienced marketers now report spending 60 to 80 percent of their time on operational execution rather than judgment.¹ Three pressures are making that arrangement harder to defend with each quarter that passes.

Retail media is fragmenting faster than measurement can follow

Brands are engaged with an average of six retail media networks today and expect that number to reach 11 by the end of 2026.⁵ Meanwhile only 15 percent of retail media advertisers report strong confidence in their measurement, and when asked what blocks them from measuring incrementality, 56 percent point to limited analytics and resources rather than anything conceptual.⁵ You're being asked to nearly double the number of networks you buy while your ability to tell which ones worked is already the weakest link. Adding headcount doesn't fix that, because the work is reconciliation, and reconciliation is exactly what people are often worst at and slowest doing.



01 • THE TIMING PROBLEM, CONTINUED

Discovery is being rewritten around the best answer, not the biggest brand

NIQ projects that agentic shoppers could drive between \$190 billion and \$385 billion in US ecommerce spending by 2030, or 10 to 20 percent of the market, and frames the shift bluntly: AI agents optimize for the best answer, not the biggest brand.⁶ Nearly half of agentic search results draw on creator ecosystems such as YouTube, TikTok, and Reddit rather than on brand-owned properties.⁶

If you've spent decades buying reach and defending share of voice, this is the part you need to pay particular attention to. The new gatekeeper doesn't recognize your brand as a reason to recommend you. It recognizes structured, well-evidenced answers to a shopper's stated goal, and hands the recommendation to whoever supplies those.

“


AI agents optimize for the best answer, not the biggest brand.”

NIQ, MARCH 2026

AGENTIC SHOPPERS' PROJECTED SHARE OF US ECOMMERCE SPENDING, 2030

\$190–385B

in US ecommerce spending



10–20% of the market (low to high estimate)

100%

Source: NIQ, Agentic Commerce and AI in CPG.⁶

The category read arrives after the decision window closes

Syndicated category data does what it was designed to do, on a drop schedule. The problem is that the buyer meeting, the markdown call, and the promotional commitment don't wait for the drop, so category managers walk in with last month's export and defend a position they can't yet see the evidence for.

Our recommendation here is specific: keep your syndicated provider and stop asking it to be an early-warning system. It was never built to be one. What you're missing isn't better historical data, it's strong analysis and recommendations provided against the signals and trends based on movement between the drops.

WHERE THE DECISION WINDOW SITS**WEEK 1****The category moves**

A competitor promotes, an assortment shifts, a trend starts.

WEEK 2–3**The decisions happen**

Buyer meeting, markdown call, promotional commitment.

NEXT DROP**The read arrives**

Accurate, and after the window has closed.

02 • DEFINING THE AGENT

What an agent is, and what it is not

The word agent now gets applied to almost anything with a text box, which has made it useless as a buying signal. The distinction that matters is architectural, and we like to say that it's best understood as four levels rather than a yes or no.

LEVEL 1	Traditional automation is deterministic. It executes if-then logic reliably and auditably, and it breaks the moment conditions change in a way nobody anticipated.
LEVEL 2	Machine learning models find patterns in historical data and predict well, but their scope is narrow and a human still has to act on the recommendation.
LEVEL 3	Large language models and chatbots understand language and shorten the path to a first draft, but they're fundamentally reactive. They respond to a prompt and can't take action in your systems on their own.
LEVEL 4 • AGENTS	Agents combine the language understanding of a model with memory, tool access, multi-step planning, and the ability to execute across systems. The difference isn't incremental. It's architectural. ¹

“ A chatbot generates a response. **An agent generates a plan, executes it, evaluates the results, and adapts.**”

Five components make that work: **the model**, which supplies reasoning; **the prompt**, which encodes your objectives and brand rules; **memory**, which accumulates business context so the system improves the longer you run it; **tools**, the systems and APIs the agent acts through, which give it hands and not just a brain; and **history**, the record of what it decided and why, which serves as both a learning mechanism and an audit trail.¹

FIVE COMPONENTS OF AN AGENT

Model Supplies reasoning	Prompt Encodes your objectives and brand rules	Memory Accumulates business context	Tools The systems the agent acts through	History What it decided and why
------------------------------------	--	---	--	---

02 • DEFINING THE AGENT, CONTINUED

Three levels of capability, in CPG terms

Capability then arrives in three levels, and here's how that translates for CPG companies:

01

SENSE

On-demand agents

Activate when you ask and complete a defined task. For a category team, that can mean assembling a competitive and category read across dozens of public sources in minutes rather than across two days of tab-switching.

02

DECIDE / RECOMMEND

Continuous agents

Monitor in the background and act on thresholds. That can mean watching a competitor's promotional pattern, a shift in a retailer's assortment, or the leading indicators of a category move, and raising it while you can still respond.

03

ACT

Autonomous agents

Pursue an outcome with minimal direction, planning and reallocating media inside the guardrails you set.



Our recommendation on sequencing is to always start at the first level, prove it out, and then progress. Teams that begin at the third level tend to discover their data problems at the worst possible moment, usually in front of commercial leadership.¹

03 • THE READINESS GAP

Addressing the readiness gap

A bit of cursory research indicates that the CPG vertical is not short on AI enthusiasm. BCG surveyed 100 C-level marketing leaders, VPs, and directors at US and European CPG companies above \$1 billion in revenue and found 70 percent expect generative AI to make marketing faster and more efficient, while only 13 percent say it's in widespread use or fully integrated into workflows.⁴

EXPECTATION VS. INTEGRATION IN CPG MARKETING




Source: Boston Consulting Group, *The AI-Forward CPG Marketing Organization* (n=100).⁴

We characterize this trend not as a technology gap, but as an integration gap, and it's the clearest signal in the CPG data that the constraint sits in the plumbing rather than in the models.

An ownership question that's about to become an org chart argument

The second gap is related to who's in charge. [Our own research](#) across 225 senior leaders at large US enterprises found that 40 percent think the Chief AI Officer should own agentic marketing, and that the answer depends heavily on who you ask: 52 percent of AI leaders say the Chief AI Officer should own it, against 25 percent of marketing executives.²


WHO SAYS THE CHIEF AI OFFICER SHOULD OWN AGENTIC MARKETING



Source: Kana, *The Agentic Divide* (n=225).²

03 • THE READINESS GAP, CONTINUED

In CPG, BCG found that nearly 70 percent of leaders now have a Chief Growth Officer or equivalent whose remit extends past marketing, nearly 60 percent are centralizing functions, and nearly 50 percent are simultaneously localizing execution.⁴




Source: Boston Consulting Group.⁴

That disagreement won't remain academic. It's going to become a live org chart argument that will decide whether your agentic program has a budget owner or a committee. The same research ranked vendor complexity among the top obstacles to agentic marketing, at 3.29 on a five-point scale,² which points at the obvious trap:

“ The answer to fragmentation can't be another point solution that fragments it further.”

Epsilon's 2026 benchmark study of more than 250 marketing decision-makers across CPG, retail, financial services, travel, and restaurants fills in the rest. Every marketer surveyed is using AI, but 71 percent are using it for productivity and only 9 percent for revenue generation, data quality is the most-cited challenge at 45 percent, and confidence tracks seniority rather than evidence: 73 percent of C-level respondents rate their AI tools extremely valuable, against 25 percent of senior managers.⁷

HOW MARKETERS ARE USING AI TODAY



Source: Epsilon, 2026 benchmark study (n=250+).⁷

Put those together and the failure mode is easy to name. CPG organizations are adopting AI quickly, using it mostly to make existing work faster, assessing their own foundations from the altitude furthest from the work, and arguing about ownership while the pilots stall. Our recommendation is to settle the ownership questions and the data-access questions when you scope the pilot, not after it has begun to slump into a state of disappointment.

04 • ENTERPRISE GRADE

What enterprise grade has to mean in CPG

Most enterprise AI failures aren't model failures. They come from missing context, weak governance, shallow evaluation, or fragile integrations.³

In this section we make recommendations on requirements, and provide some background on what each requirement provides your company and team in terms of protection.

Integration that writes as well as reads

A system that only reads is a reporting tool with better manners. Agentic marketing requires both data in and data out: the ability to update records, trigger workflows, and create tasks in the systems where your work actually happens.³ It also has to ingest the system of record, meaning the data itself, and the system of engagement, meaning the activity and context that produced it.³

For a CPG brand that generally means your data warehouse, CRM, marketing automation platform, and every retail media network you buy, connected without migrating data.

Knowledge architecture rather than a vector database

The layer that makes agents useful over time is a structure that maps entities, relationships, and context, and compounds value as the data feeding it grows. It has to capture domain-specific entity relationships rather than generic ontologies, temporal context so it knows how relationships have changed, decision provenance, and failure patterns so that unsuccessful outcomes teach the system something.³

“

Agentic marketing requires both data in and data out: the ability to update records, trigger workflows, and create tasks in the systems where your work actually happens.”

Generic foundation models know the world. **Your knowledge graph knows your business.**³

IN CPG, YOUR BUSINESS IS

Category hierarchy

Trade calendar

Retailer relationships

Pack architecture

Promotional economics

No general model has ever seen any of it.

04 • ENTERPRISE GRADE, CONTINUED

Governance enforced before the action

Access control has to mirror your organization: role-based permissions, data-level visibility aligned to human permissions, distinct authorization for reading versus deleting, and audit inheritance so an agent's action carries the identity of the person it acted for.³ Here's the practical version: if your team touches retailer-confidential or shopper-level data, ask any vendor a direct question: does the system check permissions before the agent acts, or does it record what happened afterward? Only the first one keeps you compliant, and only the first one will survive a retailer's audit.

An audit trail that survives a question

"The agent made a mistake" is not an acceptable incident report.³ A complete trail should record every decision with its reasoning, every interaction with its inputs and tools, every response including the alternatives rejected, and every failure with its recovery path, held immutably and searchably.³ This angle is incredibly important because it means the difference between a category read your commercial team will take into a buyer meeting, and one they'll quietly rebuild in Excel first.

WHAT A COMPLETE TRAIL RECORDS

Every decision

with its reasoning

Every interaction

with its inputs and tools

Every response

including alternatives
rejected

Every failure

with its recovery path

Evaluation as monitoring, not as a launch gate

Evaluation isn't pre-deployment testing. It should be considered continuous monitoring.³ Factual accuracy, tool-use correctness, compliance checks, and cost budgets all have to run against production interactions, with degradation raising an alert. CPG seasonality guarantees that a model validated in a quiet month will meet conditions in peak it has never seen, so don't give this one the short shrift.

Bounded risk instead of promised perfection

Any vendor promising zero hallucinations is committing to something that is frankly unachievable, and it's the wrong goal anyway. The goal is bounded risk: retrieval grounded in verified sources, generation constrained to known entities, citation verification, consistency checks against the system of record, confidence thresholds for autonomous action, and human review for high-stakes decisions.³ Alongside that sits a requirement we'd encourage you to test in the demo itself: determinism. A system where the same question produces different answers across successive calls can't support a compliance workflow or a buyer meeting.³ Ask the same question three times and watch what happens.

04 • ENTERPRISE GRADE, CONTINUED

AI proposes, you decide

The posture that holds up in practice is augmented intelligence. Recommendations appear in the tools your team already uses, humans review the output, and autonomy expands as trust is earned, not as a default setting. And recommendations have to activate, because a recommendation that stops at a dashboard hasn't been delivered.³

“

Autonomy expands as trust is earned, not as a default setting.”

THE SIX TESTS OF ENTERPRISE GRADE, AT A GLANCE

01

Integration that writes as well as reads

Data in and out across warehouse, CRM, marketing automation and every retail media network.

02

Knowledge architecture, not a vector database

Your category hierarchy, trade calendar and promotional economics, compounding over time.

03

Governance enforced before the action

Permissions checked up front, so retailer-confidential data stays protected.

04

An audit trail that survives a question

Every decision, input, rejected alternative and failure, recorded and searchable.

05

Evaluation as continuous monitoring

Accuracy, tool use, compliance and cost checked in production, through peak season.

06

Bounded risk, determinism, humans who decide

Same question, same answer, and autonomy that expands only as trust is earned.

05 • BUILD-WITH

The build is not a brand marketer's job

CPG marketing leaders are usually offered two options for procuring new software. Both are versions of the same arrangement enterprise software has run for thirty years, in which the vendor controls the roadmap and the customer absorbs the implementation.⁸

OPTION 1**Build**

Expensive and brittle, and never finishes.

OPTION 2**Buy**

Fits generically, then costs a lot of time, money, and resources to integrate.

THE THIRD OPTION**Build-With**

Shaped to your workflow, at the speed and cost of buying.

The reason packaged martech disappoints CPG specifically is that the work only looks recognizable across companies in its outline. Inside that outline it fragments into dozens of functions that look meaningfully different at every organization. The packaged stack handles the outline. The work lives in the fragments.⁸ Your category hierarchy isn't your competitor's. Your trade terms aren't either. A template that fits both fits neither.

“

The packaged stack handles the outline. The work lives in the fragments.”

There's a third option, which **we call Build-With**: software shaped to your workflow, configured to your data, governed by your policies, and tuned to your tolerances, at the speed and cost of buying, with the durability of a vendor relationship.⁸

WHAT BUILD-WITH IS NOT

Not low-code	which shipped the engineering problem back to you in a friendlier wrapper.
---------------------	--

Not consulting	which produces one-off systems that rot.
-----------------------	--

Not a platform	which ships the integration burden to you. ⁸
-----------------------	---

We deliver a complete application, not a kit. Your control is real but bounded: you tune outcomes, set policy, and instrument results. You don't extend, fork, or maintain.⁸

05 • BUILD-WITH, CONTINUED

Who holds the maintenance burden three years from now?


That distinction is really important because it decides who holds the maintenance burden three years from now. Many AI-native vendors chose to expose agent studios, builder toolkits, and SDKs. We deliberately didn't, because every agent you build in a studio is a divergent artifact you now own.⁸ Somebody has to update it when the model changes, the data evolves, the regulation shifts, or the platform updates. On a brand marketing team, that somebody frankly doesn't exist. We prefer that the technical burden, what we call the technical tax that every enterprise AI deployment carries, belong to us rather than to you.⁸

WHO CARRIES THE TECHNICAL TAX

When this changes	Agent studio	Build-With
The model	Your team updates every agent	Kana absorbs it
The data	Your team updates every agent	Kana absorbs it
A regulation	Your team updates every agent	Kana absorbs it
The platform	Your team updates every agent	Kana absorbs it

Shared infrastructure no brand should solve alone

Underneath the applications sits shared infrastructure no individual brand should be solving alone: governance and guardrails, LLM routing and cost optimization, access control, nondeterminism management, knowledge graph compilation, evaluation and observability, and data connectors into and out of your existing systems.⁸



“ Each is a problem every enterprise AI deployment must solve. None is a problem your marketing organization should have to solve.”⁸

06 • IN PRACTICE

What agentic marketing means when applied to CPG

So how does agentic marketing actually get applied in a CPG context? What are the low-hanging fruit use cases where you'd feel the impact soonest? Here are four, in the order we'd recommend approaching them.

01

Marketing Intelligence

02

Category Intelligence

03

Audience Builder

04

Omni-Channel Media
Planner**USE CASE 01 • MARKETING INTELLIGENCE****See what your ad spend actually sold**

Start here. [Kana Marketing Intelligence](#) unifies your campaign and CRM data into one view you can query in plain language, which turns the performance conversation from a negotiation between platform-reported numbers into an evidence-backed answer. With retail media heading toward 11 networks per brand and only 15 percent of advertisers confident in their measurement, the team that can answer “what did this actually sell” with one number becomes the team commercial leadership believes. We recommend it first because it's also the connection that makes everything after it more accurate.

USE CASE 02 • CATEGORY INTELLIGENCE**Catch the trend before it hits POS**

[Kana Category Intelligence](#) researches category questions across hundreds of curated public sources, cross-referencing analyst reports, trade press, and search trend data automatically, and answers in seconds with citations and dates attached. To be clear about what this is not: it doesn't replace your Nielsen or Circana subscription. Category managers can then walk into the buyer meeting with this week's signals instead of last month's export, and the weekly category pack assembles itself as a byproduct.

SYNDICATED DATA

Tells you what sold.

CATEGORY INTELLIGENCE

Tells you what's moving.

06 • IN PRACTICE, CONTINUED

USE CASE 03 • AUDIENCE BUILDER

Run the brief you actually wrote

[Kana Audience Builder](#) turns a plain-language brief into a live, activated segment without SQL and without an engineering ticket, and keeps every segment traceable back to the brief that requested it. That traceability is the part brand teams underrate. Targeting drift is rarely the result of a single bad decision; it's actually the byproduct of a dozen small translations between what the brief asked for and what the platform could execute. For new SKUs and markets where you don't have first-party data yet, synthetic audiences built from public signals give you a defensible starting point on day one.

USE CASE 04 • OMNI-CHANNEL MEDIA PLANNER

Stop optimizing with last year's channel mix

Most CPG media plans are inherited, and the platforms reporting performance are the same platforms asking for more budget. The [Kana Omni-Channel Media Planner](#) models saturation curves from your own campaign history rather than from industry benchmarks, and lets you plan at minus 20 percent, baseline, and plus 20 percent so you walk into the budget conversation with a model instead of a guess. You can show which channels have hit diminishing returns before you commit the money rather than after.

-20%

budget scenario

Baseline

current plan

+20%

budget scenario

All of this on top of your stack, not instead of it

None of this work to adopt agentic marketing requires data migration. We connect by API and sit on top of Databricks, Salesforce, Marketo, your warehouse, and your retail media platforms, with zero-copy access and no rip-and-replace. We're actually here to help you get more return on your prior martech and adtech investments, not to add another line item competing with them. Connection requires from days to weeks, handled by our solutions team, and most CPG teams see value inside 30 days by connecting their top three to five data sources and configuring Marketing Intelligence first, then layering in Category Intelligence and Audience Builder once those connections are stable.

FROM CONNECTION TO VALUE

Days—weeks

to connect, handled by our solutions team

3–5

top data sources connected to start

30 days

for most CPG teams to see value

07 • THE ENGAGEMENT

What a build-with agentic marketing engagement requires

Build-With is a commitment rather than a capability, so it's worth being explicit about what it asks of you. We embed Forward Deployed Engineers, senior technical specialists working inside your environment for the duration of onboarding.⁸ In return, five things are asked of you:

WHAT WE ASK OF YOU

- 01 Read access to the relevant systems** (which, if not granted efficiently, is the single most common source of delay in onboarding)
- 02 A clear account** of how audiences are defined, campaigns approved, and category and trade decisions made
- 03 Specific and measurable success criteria**
- 04 The governance and compliance constraints that apply**, including anything your retailer agreements impose
- 05 An executive sponsor plus a working lead** who will genuinely engage⁸

WHAT YOU RECEIVE IN EXCHANGE

A production-grade application

Deployed, rather than a prototype or a sandbox.

A knowledge graph

That captures your institutional knowledge and compounds with continued use.

Governance training

Sufficient to own the outcomes.

Us as the ongoing owner of maintenance

As models improve, platforms change, and regulations shift.⁸

CPG marketing has always had judgment. What it has lacked is a system worthy of it.⁸ The translation layer between a category manager's judgment and the systems that execute on it has been staffed by people, governed by ticket queues, and measured in weeks. Collapsing that layer is the point. We act on what's happening now.⁸

08 • CONCLUSION

One system, finally ahead of the category

CPG marketing teams don't need another tool in the stack. You need the layer that makes the stack and the data behave as one system, and that's what Kana's marketing operating layer delivers. We connect to the warehouse, CRM, marketing automation platform, and retail media networks you already run, with zero-copy access and no migration, and we run the full lifecycle as a single loop rather than a diagram: reading the category as it moves, translating a brief into a traceable audience, modeling where the next media dollar still works, and answering what your spend actually sold with one number drawn from your own data.

Because your agentic marketing applications are built with you rather than handed over as a kit, the knowledge graph underneath them compounds with every interaction instead of decaying, and the technical tax of model updates, platform changes, and regulatory shifts stays with us rather than landing on a marketing team that was never staffed to pay it.

For a CPG organization facing a retail media footprint on its way to eleven networks, a discovery layer that now rewards the best answer over the biggest brand, and a commercial leadership team that wants evidence rather than platform self-grades, that combination of speed, governance, and durability is the difference between piloting agentic marketing and operating it.

Ready to get started?

We'd love to book 30 minutes with you to discuss what this would look like on your own data.

Book 30 minutes at kana.ai →



ENDNOTES

Sources

All third-party sources cited below were published within the trailing 24 months.

- 1 Kana, *Understanding AI Agents for Marketing*.
- 2 Kana, *The Agentic Divide*, June 2026. Survey of 225 senior leaders at US enterprises with \$250M or more in revenue or 3,000 or more employees, fielded May 2026.
- 3 Kana, *Evaluation Guide: Is Your AI Enterprise Grade?*
- 4 Boston Consulting Group, *The AI-Forward CPG Marketing Organization*, 12 February 2026. Survey of 100 C-level marketing leaders, vice presidents, and directors at US and European companies with \$1 billion or more in annual revenue.
- 5 Skai, *The 2026 State of Retail Media Fragmentation: 7 Challenges Brands Must Solve to Scale*, 12 February 2026. Survey of 166 retail media advertisers conducted with Stratably.
- 6 NIQ, *Agentic Commerce and AI in CPG: The New Growth Frontier*, 2 March 2026.
- 7 Epsilon, *2026 benchmark study: Marketing's AI inflection point*, 10 August 2026. Survey of more than 250 marketing decision-makers across retail, CPG, financial services, travel, and restaurants.
- 8 Kana, *The Build Is Not Your Job*, June 2026.

CPG capability descriptions and integration details are drawn from Kana's published CPG materials.

ABOUT KANA

Kana is an agentic marketing platform that builds custom AI applications matched to your business goals and architecture, offering the precision of a custom build with enterprise-grade scale and security. Kana synthesizes disparate data into one system that learns, decides, and acts through the intelligent orchestration of AI agents, empowering marketers to make smarter decisions that deliver faster growth and scale.

Founded in 2025, backed by Mayfield, and headquartered in San Francisco, California.