

Finally Ahead of the Curve

An exploration of the impact that agentic marketing will have on retail and eCommerce, and what it demands of marketing leadership moving forward.

CONTENTS

Inside this paper

—	Executive summary	3
01	Retail'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 retail	10
05	The build is not a retail marketer's job	13
06	What agentic marketing means when applied to retail	15
07	What a Build-With engagement requires	17
08	Conclusion	18
—	Endnotes	19

EXECUTIVE SUMMARY

Late decisions are the most expensive kind

Retail has a timing problem that no amount of incremental tooling has been able to resolve.

A customer lapses in March and the retention report names them in May. A category shifts on a Tuesday and the trade meeting hears about it the following Monday. A promo window opens and closes while the audience request sits in a queue. Each of those gaps is a margin decision made late, and late decisions are the most expensive kind of decision that retail marketing teams make.

Early testing indicates that agentic AI will be able to successfully close that gap. It's not a faster dashboard and it's not a better chatbot. It's actually 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. Agentic marketing applies that principle to your marketing and go-to-market strategy. Retail is where that architecture actually matters most, because retail runs on narrow windows and thin margins, and because the discovery layer that feeds retail demand is being rebuilt around AI right now.

This white paper has been 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 retail organizations really are to adopt agentic AI
- 03 What enterprise grade has to mean when the data in question is your register, your loyalty file, and your customer list
- 04 How agentic marketing applies to the specific decisions retail teams make every week

This paper was written for retail marketing leaders who have been sold transformation before and want to understand what is structurally different this time around. We hope that it will help the reader gain more trust in agentic marketing and its implications for the next era of retail marketing. Let's dive in.

01 • THE TIMING PROBLEM

Retail's timing problem is an architecture problem

Often the instinct when marketing feels slow is to buy something that promises to speed it up. That instinct largely contributed to the current problem that marketers in every industry (not just retail) are now facing. In 2011 there were roughly 150 martech solutions. Today there are more than 15,000.¹ In that time, the majority of retail teams have accumulated at least a commerce platform, an ESP, a loyalty platform, a CDP, a POS system, three or more ad platforms, and an analytics layer that reconciles none of them cleanly.

THE STACK GREW. THE CONNECTIVE TISSUE DIDN'T.**~150**martech solutions on the
market in 2011**15,000+**martech solutions on the
market today**60–80%**of practitioner time spent on
operational execution, not
judgmentSource: Kana, Understanding AI Agents for Marketing.¹

Every tool in that stack works on its own. What was never built is the connective tissue that makes them work as a system, and so the connective tissue became the marketer. Experienced practitioners now spend the majority of their time on operational execution rather than judgment, somewhere in the range of 60 to 80 percent.¹

As stated above, that is an architecture problem, not a tooling problem, and it shows up in retail as a persistent lag between when something becomes true and when the organization can or will act upon it. The lag has a cost that rarely appears as a line item. It appears as a win-back discount offered to a shopper who would have returned at full price. It appears as markdown taken a week later than it should have been. It appears as a campaign that launched after the promo window closed.

01 • THE TIMING PROBLEM, CONTINUED

Two forces making the lag less survivable

The first is that most retail revenue still closes where the digital measurement stack cannot see it. Ecommerce accounted for 17.1 per cent of total US retail sales in the second quarter of 2026, which means roughly 83 per cent of retail sales closed somewhere other than a website.⁵ Every ad platform still counts conversions its own way, and the CFO is frequently asked to accept a number that no two (or more) systems agree on.

WHERE US RETAIL SALES CLOSE, Q2 2026




Source: US Census Bureau via FRED.⁵

The second force is that retail discovery is being rebuilt around AI faster than most retail measurement has adjusted. Traffic from AI sources to US retail sites grew 393 percent year over year in the first quarter of 2026.⁷ Measured from October 2024, AI-referred traffic to retail sites has grown more than fourteenfold, and by May 2026 that traffic was converting 54 percent better than traffic from non-AI sources, with visitors spending 53 percent more time on site and browsing 23 percent more pages per visit.⁸ A year earlier the same channel converted worse than everything else.⁷

A channel that inverts from worst to best in twelve months is not something a quarterly planning cycle can necessarily track.

AI-REFERRED RETAIL TRAFFIC VS. NON-AI SOURCES, MAY 2026



Source: Digital Commerce 360, citing Adobe Analytics.⁸

02 • DEFINING THE AGENT

What an agent is, and what it is not

The word agent has been applied to almost everything, which has made it useless as a buying signal. The distinction between what an agent and automation is best understood as four levels of capability, rather than a binary.

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 are fundamentally reactive: they respond to a prompt and cannot 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 is not incremental. It is architectural. A chatbot generates a response. An agent generates a plan, executes it, evaluates the results, and adapts. ¹

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

02 • DEFINING THE AGENT, CONTINUED

Three levels of capability, in retail terms

Capability then arrives in three levels, and the retail translation of each is concrete.

**01**

START HERE

On-demand agents

Activate when asked and complete a defined task. In retail, that can mean consolidating performance across three ad platforms, the POS, and the loyalty file into one view in minutes rather than an afternoon of exports.

**02**

THEN

Continuous agents

Monitor in the background and act on thresholds. That is detecting a cost-per-acquisition increase and pausing the ad sets driving it before the weekly review, or flagging a cohort whose purchase frequency has started to decay.

**03**

AS TRUST BUILDS

Autonomous agents

Pursue an outcome with minimal direction, designing, deploying, monitoring, and iterating a retention program inside the guardrails a human set.

And the sequencing matters. Organizations should begin at the first level, progress to the second, and expand into the third as governance and confidence build.¹ Retail teams that invert that order may find that they discover their data problems at the worst possible moment.

03 • THE READINESS GAP


Addressing the readiness gap

The market has already decided that agentic marketing is happening. Gartner expects 40 percent of enterprise applications to include task-specific AI agents by the end of 2026, up from fewer than 5 percent in 2025, and Forrester expects one quarter of shoppers to use specialty retail chatbots during the same year.⁶ Kana's own research across 225 senior leaders at large US enterprises found that 82 per cent expect agents to run a third or more of their routine decisions within two years, and 46 per cent expect agents to handle a majority of routine decisions inside 24 months.² Whatever one makes of the forecasts, the overarching assumption is now widely shared.

THE ASSUMPTION IS WIDELY SHARED

40%

of enterprise applications expected to include task-specific AI agents by the end of 2026, up from fewer than 5% in 2025 (dashed line)⁶




82%

of senior leaders expect agents to run a third or more of their routine decisions within two years²



46%

expect agents to handle a majority of routine decisions inside 24 months²



Sources: National Retail Federation, citing Gartner; Kana, The Agentic Divide (n=225).

What is less shared is an accurate view of readiness. In the same research, 86 per cent of leaders called their data infrastructure ready for agentic marketing.² Set that beside Epsilon's 2026 benchmark study of more than 250 marketing decision-makers across retail, CPG, financial services, travel, and restaurants which found that data quality (described as incomplete, inconsistent, or unreliable data) to be the single most-cited challenge, named by 45 per cent of marketers.⁹ Both numbers cannot be describing the same condition.

03 • THE READINESS GAP, CONTINUED


The Epsilon study also found that every marketer surveyed is using AI. 71 per cent reported using it for productivity and efficiency, and only 9 per cent are using it for revenue generation.⁹ It also found that confidence tracks seniority rather than evidence: 73 per cent of C-level respondents considered their implemented AI tools extremely valuable, against 25 per cent of senior managers, and half of all marketers rated their organization extremely mature with AI, up from 35 per cent a year earlier.⁹

READINESS, AS REPORTED VS. AS EXPERIENCED


Leaders who call their data infrastructure ready² **86%**



Marketers naming data quality their top challenge⁹ **45%**



Marketers using AI for revenue generation⁹ **9%**



Sources: Kana, The Agentic Divide; Epsilon, 2026 benchmark study.

Read together, these findings describe a specific but avoidable failure mode. Retail organizations are adopting AI quickly, using it mostly to make existing work faster, assessing their own foundations generously, and forming that assessment furthest from the work. That is how a pilot becomes a stalled program. But it is also why the question a retail CMO should be asking is not which AI features a vendor offers, but what the vendor is prepared to own. We'll explore that more in detail in the next section.

04 • ENTERPRISE GRADE

What enterprise grade has to mean in retail

Most enterprise AI failures are not model failures. They are the result of missing context, weak governance, shallow evaluation, or fragile integrations.³

In retail those four failure modes have specific shapes, and a serious evaluation should test for them directly.

Integration that writes as well as reads

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

For a retailer that means commerce platform, POS, loyalty, ESP, ad platforms, and warehouse, connected without a migration.

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 so the reasoning behind an action is recoverable, and failure patterns so unsuccessful outcomes teach the system something.³

The principle is short: generic foundation models know the world, and your knowledge graph knows your business.³ In retail, that business is your category structure, your promotional calendar, your store footprint, and your loyalty economics, none of which a general model has ever seen.



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

04 • ENTERPRISE GRADE, CONTINUED

Governance enforced before the action

Access control has to mirror the 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 on whose behalf it acted.³ Enforcement has to happen before the action, not in a report afterward. It bears emphasizing this point: retail teams handling loyalty and payment-adjacent data should treat pre-action enforcement as a threshold requirement rather than a feature.



Retail teams handling loyalty and payment-adjacent data should treat pre-action enforcement as a threshold requirement rather than a feature.”

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, outputs, and tools used, every response including the alternatives rejected, and every failure with its classification and recovery path, held immutably and searchably.³ That trail serves compliance, debugging, and training at once, and it is the difference between a system you will trust with a markdown call and one you'll quietly route around.

Evaluation as monitoring, not as a launch gate

Evaluation is not pre-deployment testing. It should be considered continuous monitoring.³ Factual accuracy against internal sources, tool-use correctness, safety and compliance checks, and latency and cost budgets all have to run against production interactions, with degradation raising an alert. Retail's seasonality guarantees that a model validated in April will meet conditions in November it has never seen, and it must be trained and built to anticipate that seasonality.

Bounded risk instead of promised perfection

Any vendor claiming zero hallucinations is describing something unachievable, and the wrong goal in any case. The goal must be 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, human review for high-stakes decisions, and graceful degradation when uncertainty is high.³

Alongside that sits a requirement retail leaders tend to underrate: determinism. A system in which the same question produces different answers across successive calls cannot support a compliance workflow, a validation process, or a trade meeting.³

04 • ENTERPRISE GRADE, CONTINUED

AI proposes, humans decide

The posture that holds up in practice is augmented intelligence. Recommendations appear in the tools teams already use, review interfaces exist for approval, handoff protocols are clear, and feedback loops improve what comes next. Autonomy expands as trust is earned and evaluations prove reliable, not as a default setting. And recommendations have to activate: a recommendation that stops at a dashboard has not been delivered.³

THE SIX TESTS OF ENTERPRISE GRADE, AT A GLANCE

01

Integration that writes as well as reads

Data in and data out: update records, trigger workflows and create tasks across commerce, POS, loyalty, ESP and ads.

02

Knowledge architecture, not a vector database

Entities, relationships, time, provenance and failure patterns, compounding as your data grows.

03

Governance enforced before the action

Role-based permissions and audit inheritance, enforced up front rather than reported afterward.

04

An audit trail that survives a question

Every decision, input, rejected alternative and failure recorded immutably and searchably.

05

Evaluation as continuous monitoring

Accuracy, tool use, safety, latency and cost checked against production, with alerts on degradation.

06

Bounded risk, determinism, humans who decide

Grounded, consistent output, and autonomy that expands only as trust is earned.

05 • BUILD-WITH

The build is not a retail marketer's job

Historically, retail marketing leaders have been offered two options for procuring new software solutions:

- **Build**, which is expensive and brittle and never finishes.
- **Buy**, which fits generically and then costs a lot of time, money, and resources to integrate.

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.⁴ The reason packaged martech disappoints retail, specifically, is that the work looks recognizable across companies only in its outline. Inside that outline it fragments into dozens of functions that look significantly different within every organization. The packaged stack handles the outline, but the work lives in the fragments.⁴



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

There is a third option, which Kana calls Build-With: software shaped to a retailer's workflow, configured to their data, governed by their policies, and tuned to their tolerances, at the speed and cost of buying, with the durability of a vendor relationship.⁴

At Kana, we frequently prefer to characterize our definition of "Build-With" by dint of what it is not, as follows:

Not low-code

Low-code shipped the engineering problem back to the customer in a friendlier wrapper.

Not consulting

Consulting produces one-off systems that rot.

Not a platform

Platforms ship the integration and orchestration burden to the customer.⁴

Kana delivers a complete application, not a kit. The customer's control is real but bounded: they tune outcomes, set policy, and instrument results. They do not extend, fork, or maintain.⁴

05 • BUILD-WITH, CONTINUED

Who holds the maintenance burden three years from now?

That distinction is consequential simply because it determines who holds the maintenance burden three years from now. Many AI-native vendors chose to expose agent studios, builder toolkits, and SDKs. Kana deliberately did not, on the reasoning that every agent a customer builds in a studio is a divergent artifact the customer now owns.⁴ Someone has to update that artifact when the model changes, the data evolves, the regulation shifts, or the platform updates. In a retail marketing team, that someone simply does not exist. We prefer for the technical burden, what Kana calls the technical tax that every enterprise AI deployment carries, to belong to Kana, not to the customer.⁴

Shared infrastructure no retailer should solve alone

Underneath the applications sits shared infrastructure that no individual retailer should be solving alone: governance and guardrails, LLM routing and cost optimization, attribute- and role-based access control, nondeterminism management that turns probabilistic output into stable, trackable, reproducible signals, knowledge graph compilation, evaluation and observability, and data connectors into and out of existing systems.⁴

Each is a problem every enterprise AI deployment must solve. None is a problem a retail marketing organization should have to solve.⁴



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

06 • IN PRACTICE

What agentic marketing means when applied to retail

How will agentic marketing be successfully applied in a retail context? What are the best low-hanging fruit use cases where the impact can be most immediately felt?

USE CASE 01 • MARKETING INTELLIGENCE

Connect digital spend to register revenue

With roughly 83 percent of retail sales closing outside the website and every ad platform grading its own homework, the CFO conversation is usually an argument about whose number is right. [Kana Marketing Intelligence](#) unifies ad, ecommerce, POS, and loyalty data with zero-copy access, so campaign spend, basket size, and customer lifetime value can be seen together, grounded in connected first-party data rather than platform-reported conversions. Plain-language questions replace hours of reconciliation, and the number that reaches the finance review survives scrutiny because it came from the retailer's own systems.

USE CASE 02 • PERSONALIZATION

Catch churn weeks before the retention report

Most first-time ecommerce customers never return, and churn typically surfaces as a lagging line in a monthly report. By publication, the at-risk customers have already lapsed, which leaves win-back as the only available move, and win-back defaults to blanket discounting that burns margin on shoppers who would have come back at full price. The [Kana Personalization](#) application can score customer churn risks ahead of the retention report, reading the CRM, loyalty, and marketing automation data already in place, and can recommend the intervention worth making with a stated reason. Retention spend stops subsidizing customers who were already returning and starts saving the ones who were not. Every recommendation goes through the team's approval.

USE CASE 03 • CATEGORY INTELLIGENCE

Walk into the Monday trade meeting with answers

Category trends, competitor moves, and promo benchmarks currently arrive as a spreadsheet marathon that ends the morning the decisions are due. [Kana Category Intelligence](#) tracks those signals continuously and answers merchant questions in seconds against a cited, dated record, which means markdown and allocation calls get made while they still change the outcome. The weekly trade pack assembles itself as a byproduct rather than as a project.

06 • IN PRACTICE, CONTINUED

USE CASE 04 • AUDIENCE BUILDER

Brief to live segment inside the promo window

An audience request that takes days is a promotional window missed, and in retail the window does not reopen. [Kana Audience Builder](#) turns a brief into a live, activated segment in minutes, in plain language, without SQL or an engineering queue, with reach estimated across Meta, Google, and TikTok before any budget commits. For new stores, new markets, or new categories where no history exists yet, synthetic audiences provide a defensible starting point on day one that refines as real signal arrives.

USE CASE 05 • OMNI-CHANNEL MEDIA PLANNER

Reallocate on modeled return, not platform self-grades

Channel mix in retail tends to be inherited rather than decided, and the platforms reporting performance are the same platforms asking for the budget. The [Kana Omni-Channel Media Planner](#) models saturation curves against a retailer's own campaign history across Meta, YouTube, The Trade Desk, and DV360, so budget follows modeled marginal return rather than last year's allocation or a platform's self-assessment.

On top of the stack, not instead of it

None of the above requires data migration. Kana connects by API and sits on top of Shopify, Salesforce, Klaviyo, the POS, and the loyalty platform, with zero-copy access and no rip-and-replace. In some ways you could say we're actually here to help you get more return on your prior martech and adtech investments. Connection typically requires from days to weeks, handled by Kana's solutions team, and most retail teams see value inside 30 days by connecting their top three to five data sources, configuring Marketing Intelligence first, then layering in Personalization and Audience Builder once those connections are stable and insights and decisions are delivering routinely and reliably.

FROM CONNECTION TO VALUE

**Days–
weeks**

to connect, handled by Kana's
solutions team

3–5

top data sources connected
to start

30 days

for most retail teams to see
value

07 • THE ENGAGEMENT

What a Build-With agentic marketing engagement requires

Build-With is a commitment rather than a capability, and it is worth being explicit about what it requires from a retail organization. Kana embeds Forward Deployed Engineers, senior technical specialists working inside the customer's environment for the duration of onboarding.⁴

WHAT KANA ASKS OF THE RETAILER

- 1 Read access to the relevant systems (which, if not efficiently granted, can be the single most common source of delay in onboarding)
- 2 A clear account of how audiences are defined, campaigns approved, and retention decisions made
- 3 Specific and measurable success criteria
- 4 The governance and compliance constraints that apply
- 5 An executive sponsor plus a working lead who will genuinely engage⁴

WHAT THE RETAILER RECEIVES

- ✓ A production-grade deployed application rather than a prototype or a sandbox
- ✓ A knowledge graph that captures institutional knowledge and compounds with continued use
- ✓ Governance training sufficient to own the outcomes
- ✓ Kana as the ongoing owner of maintenance as models improve, platforms change, and regulations shift⁴

Retail marketing has always had judgment. What it has lacked is a system worthy of it.⁴ The translation layer between a merchant's judgment and the systems that execute on it has been staffed by people, governed by ticket queues, and measured in days. Collapsing that layer is the point. Kana acts on what is happening now.⁴

08 • CONCLUSION

One system, finally ahead of the curve

Retail marketing teams don't need another tool in the stack; they need the layer that makes the stack and the data behave as one system, and that is what Kana's marketing operating layer delivers. We connect to the commerce platform, POS, loyalty file, and ad platforms you already run, with zero-copy access and no migration, and run the full lifecycle as a single loop rather than a diagram: sourcing and scoring demand, personalizing at the moment of relevance, catching churn before the report names it, and answering the trade and finance questions 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 it compounds with every interaction instead of decaying, and the technical tax of model updates, platform changes, and regulatory shifts stays with Kana rather than landing on a marketing team that was never staffed to pay it.

For a retail organization facing narrow promo windows, a discovery layer being rebuilt around AI, and a CFO who wants a number that survives scrutiny, 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 schedule a 30 minute call to discuss what this would look like on your own data.

Book a call 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 Kana, *The Build Is Not Your Job*, June 2026.
- 5 US Census Bureau, *E-Commerce Retail Sales as a Percent of Total Sales*, Q2 2026 (17.1 percent), via FRED, Federal Reserve Bank of St. Louis. fred.stlouisfed.org/series/ECOMPCTSA. Updated 18 August 2026.
- 6 National Retail Federation, *10 trends and predictions for retail in 2026*, 7 January 2026, citing Gartner (August 2025) and Forrester Research.
- 7 Adobe, *AI traffic grows but retail sites lag in AI search visibility*, 16 April 2026. Adobe Analytics data covering January to March 2026.
- 8 Digital Commerce 360, *Adobe: AI-referred traffic to retail sites doubles in a year*, 17 June 2026, citing Adobe Analytics data for May 2026.
- 9 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.

Retail capability descriptions and integration details are drawn from Kana's published retail and eCommerce materials.