

PICK UP
HERE

— A KANA WHITE PAPER

QUICK SERVICE RESTAURANTS

Nobody Eats at the National Average

How agentic AI will close the gap between when marketing conditions change and when your marketing responds

OCTOBER 2026

CONTENTS

Inside this paper

—	Executive summary	3
01	QSR's timing problem is an architecture problem	4
02	What an agent is, and what it is not	7
03	Addressing the readiness gap	9
04	What enterprise grade has to mean in QSR	11
05	The build is not a restaurant marketer's job	14
06	What agentic marketing means when applied to QSR	16
07	What a build-with agentic marketing engagement requires	18
08	Conclusion	19
—	Endnotes	20

EXECUTIVE SUMMARY

You can buy a visit. You can't buy the next one.

The value wars worked, and that's the problem. Across 18 major Quick Service Restaurant (QSR) brands, 15 had lower customer retention in May 2026 than they did in May 2024, losing an average of 1.69 percentage points. Of the 12 brands that grew transactions over the past year, 9 saw retention fall.⁴ You can buy a visit, but it turns out you can't buy the next one.

The instinct is to read that as a discounting problem. However, we'd argue it's actually a targeting problem. A national promotional calendar can't tell the difference between a market that needed the offer and one that would have converted anyway, or between an offer your kitchen can serve tonight and one it can't. So the offer goes out everywhere, the margin goes out with it, and the guest who finds out at the counter that the item is 86'd learns something about your brand that no campaign will unteach them.

The good news is that agentic AI is the first technology category that can close the gap between when conditions change in a market and when your marketing responds. 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. Agentic marketing applies that principle to your marketing and go-to-market strategy.

In this white paper we cover four topics related to how the QSR industry stands to benefit from agentic AI and how it should prepare to migrate to agentic marketing:

-
- 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 restaurant organizations really are to adopt agentic AI
 - 03 What enterprise grade has to mean when your data spans a POS, a loyalty platform, an ordering system, and thousands of stores you don't all own
 - 04 How agentic marketing applies to the specific decisions your restaurant marketing team makes every week
-

We wrote this paper for QSR 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

QSR'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 largely built the problem. In 2011 there were roughly 150 martech solutions. Today there are more than 15,000.¹ By now, most restaurant marketing teams have accumulated a POS, an ordering platform, a loyalty program, a CRM, a warehouse, agency relationships, and a co-op structure, none of which reconcile 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

Share of an experienced marketer's time spent on operational execution rather than judgment

60–80%*Source: Kana, Understanding AI Agents for Marketing.¹*

What was never built is the connective tissue that makes those tools work as a system, so the connective tissue became your team. As a result, experienced marketers now spend 60 to 80 percent of their time on operational execution rather than judgment.¹

In QSR that shows up as three specific pressures, and all three are getting worse.

01 **Traffic is falling while checks rise**, and discounting isn't fixing it


02 **Your calendar is locked** before the conditions it plans for exist

03 **Your franchisees can see it**, and they're paying for it

01 • THE TIMING PROBLEM, CONTINUED


Traffic is falling while checks rise, and discounting isn't fixing it

Overall restaurant traffic fell 2.9 percent year over year in November 2025 while the average check rose 3.8 percent, with menu price lift contributing only 0.4 percent of that.⁵ QSR visits then fell in four of the first five months of 2026. By way of example, McDonald's same-store sales growth slowed to 0.8 percent, its weakest pace in more than a year.⁶



Meanwhile the brands outperforming the category were the ones emphasizing quality and differentiated menus rather than leading with price.⁶ Our read: the category has proven that a discount buys a transaction and doesn't buy a relationship. What's missing isn't a better offer. It's the ability to know which guest, in which market, needed one.

A DISCOUNT BUYS A TRANSACTION, NOT A RELATIONSHIP



One brand where retention fell Retention held or grew

Source: Facteus transaction data, May 2024 to April 2026.⁴

01 • THE TIMING PROBLEM, CONTINUED

Your calendar is locked before the conditions it plans for exist

Day-part campaigns run on a fixed schedule, blind to weather, local events, and competitor moves. Co-op and DMA budgets get locked quarterly with no mechanism to shift spend when a market's traffic diverges from the system. Most QSR purchase decisions resolve in minutes, so a plan committed weeks out can't keep pace. Every week the comp-sales readout arrives showing variance the national calendar ignored, which makes it a record of decisions made three weeks too late.

THREE CLOCKS THAT DON'T MATCH

CO-OP AND DMA BUDGETS

Quarterly

Locked, with no way to shift spend when a market diverges

THE CAMPAIGN PLAN

Weeks out

Committed on a fixed day-part schedule, blind to weather and local events

THE GUEST'S DECISION

Minutes

Most QSR purchase decisions resolve almost on the spot

Your franchisees can see it, and they're paying for it

Franchisees contribute between 2 and 7 percent of gross sales to pooled marketing funds, and they want proof the spend moved their stores rather than the system average. When campaigns run uniformly while traffic and menu mix vary by market, they can see exactly which stores got left out. That conversation is hard to win with a national number and, frankly, impossible to win with a national number that arrived late.

This is the part of the QSR problem with no analogue in other industries. Elsewhere, marketing answers to a CMO and a CFO. In a franchised system, it also answers to thousands of owner-operators who pay into the fund and watch their own registers.



Your marketing accountability is distributed to people who hold the receipts.”

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 nearly useless as a buying signal. The distinction that matters is architectural, and we would 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 it breaks the moment conditions change in a way nobody anticipated. A time-of-day offer rule is exactly this.
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. 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 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 QSR terms

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

01

START HERE

On-demand agents

Activate when you ask and complete a defined task. That can mean building a sized, suppressed segment from your loyalty file in minutes instead of filing a two-week ticket.

02

THEN

Continuous agents

Monitor in the background and act on thresholds. That can mean watching store-level operational signals and catching an offer for an 86'd item before it sends, or noticing a market whose traffic has diverged from the system while you can still respond.

03

AS TRUST BUILDS

Autonomous agents

Pursue an outcome with minimal direction, reallocating budget toward whatever's converting inside the guardrails you set.

Start at the first level, prove it, then progress

Our recommendation on sequencing is firm: start at the first level, prove it, then progress. Each level depends on the one before it. On-demand work shows you where your data is clean and where it isn't. Continuous monitoring only earns trust once those gaps are closed. Autonomy only makes sense once the monitoring has a track record. Restaurant teams that begin at the third level tend to discover their data problems in front of a franchise advisory council, which is the worst possible room for that discovery.¹

“




Teams that begin at the third level tend to discover their data problems in front of a franchise advisory council.”

03 • THE READINESS GAP

Addressing the readiness gap

The restaurant industry is not necessarily falling short on AI adoption. In a survey of roughly 170 limited-service brands, 51 percent were already investing in AI, another 22 percent planned to start within the year, and 53 percent of those investing were applying it to marketing and personalization.³

RESTAURANTS ARE NOT SHORT ON AI ADOPTION




Source: ANA, citing the Qu State of Digital report (n≈170 limited-service brands).³

And the initial results are both real and promising. Yum Brands has sent more than 200 million AI-generated communications that were up to five times more effective than traditional approaches, and Taco Bell hit a record 41 percent digital mix in the second quarter of 2025.³

So what's holding the category back? The same survey answers it: 37 percent of restaurant brands report that siloed technology is preventing them from maximizing their AI investments.³ That's not a model problem or a talent problem. It's a plumbing problem, and it's the clearest signal in the restaurant data that the constraint sits between your systems rather than inside any one of them.


THE CONSTRAINT SITS BETWEEN YOUR SYSTEMS



03 • THE READINESS GAP, CONTINUED


Team skills and acumen compound a growing fear of falling to the competition

Two findings from [our own research](#) across 225 senior leaders at large US enterprises sharpen the picture. First, 69 percent are concerned about competitive disadvantage from moving too slowly, which in a category where a competitor's promotional move lands in your market the same week is not something worth shrugging off as an abstract concern.² Second, when we asked what actually obstructs agentic marketing, team skills and acumen ranked highest of all barriers at 3.44 on a five-point scale, above data quality, budget, and executive alignment.² For a marketing organization running thousands of restaurants with a team sized for a brand rather than a footprint, that should land hard. The binding constraint is people-hours, and the question for many restaurant marketing teams is whether you solve that by hiring or by putting in the right architecture.



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

Epsilon's 2026 benchmark study of more than 250 marketing decision-makers across restaurants, retail, CPG, financial services, and travel 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



Our recommendation is blunt. Don't scope an agentic pilot around a campaign. Scope it around a connection, because the 37 percent figure saying their data is stuck in silos says the connection between systems is what's actually broken, and a pilot that succeeds on top of siloed data has proven nothing you can scale.

04 • ENTERPRISE GRADE

What enterprise grade has to mean in QSR

Most enterprise AI failures aren't model failures. They come from missing context, weak governance, shallow evaluation, or fragile integrations.⁸

Here's what we think you should require, and what each requirement protects you from.

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.⁸ 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.⁸ For a restaurant brand that means your POS, ordering platform, loyalty program, CRM, warehouse, and ad platforms, 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, and **failure patterns so unsuccessful outcomes teach the system something**.⁸

The principle is short: generic foundation models know the world, and your knowledge graph knows your business.⁸ In QSR, that business is your day-part structure, your menu architecture, your franchise map, your co-op boundaries, and your loyalty economics. No general model has ever seen any of it.

“

Generic foundation models know the world, and your knowledge graph knows your business.”

WHAT YOUR KNOWLEDGE GRAPH KNOWS THAT NO GENERAL MODEL DOES

Day-part structure

Menu architecture

Franchise map

Co-op boundaries

Loyalty economics

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.⁸

Here's the practical version, and it's sharper in QSR than almost anywhere. A franchisee's store data is not corporate's to use however it likes. 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 survives a franchise agreement.

“

A franchisee's store data is not corporate's to use however it likes.”

An audit trail that survives a question

“The agent made a mistake” is not an acceptable incident report.⁸ 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.⁸ In a franchised system that trail has a second job: it's how you answer a franchisee who wants to know why their market's budget moved. “The model decided” loses that meeting. A dated, reasoned record wins it.

Evaluation as monitoring, not as a launch gate

Evaluation isn't pre-deployment testing. It should be considered continuous monitoring.⁸ Factual accuracy, tool-use correctness, compliance checks, and cost budgets all have to run against production interactions, with degradation raising an alert. A limited-time offer cycle guarantees that a system validated on last quarter's menu will meet conditions next quarter it has never seen.

Bounded risk instead of promised perfection

Any vendor promising zero hallucinations is describing something unachievable, and the wrong goal anyway. The goal is bounded risk: retrieval grounded in verified sources, generation constrained to known entities, consistency checks against the system of record, confidence thresholds for autonomous action, and human review for high-stakes decisions.⁸ 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 co-op reporting obligation.⁸ 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, and autonomy expands as trust is earned, not as a default setting. And recommendations have to activate. A recommendation that stops at a dashboard hasn't been delivered.⁸

THE SIX TESTS OF ENTERPRISE GRADE, AT A GLANCE

01

Integration that writes as well as reads

Data in and data out across POS, ordering, loyalty, CRM, warehouse and ad platforms.

02

Knowledge architecture, not a vector database

Day-parts, menus, franchise map, co-op boundaries and loyalty economics, compounding over time.

03

Governance enforced before the action

Permissions checked before an agent acts, so store-level data stays where franchise agreements put it.

04

An audit trail that survives a question

A dated, reasoned record of why a market's budget moved.

05

Evaluation as continuous monitoring

Accuracy, tool use, compliance and cost checked in production, through every LTO cycle.

06

Bounded risk, determinism, you decide

The same question gets the same answer, and autonomy expands only as trust is earned.

05 • BUILD-WITH

The build is not a restaurant marketer's job

QSR marketing leaders are usually offered two options for procuring new software.

01

Build

Expensive and brittle, and it never finishes.

02

Buy

Fits generically, 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.⁹ The reason packaged martech disappoints QSR 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.⁹ Your franchise structure isn't your competitor's. Neither is your day-part mix or your co-op math. A template that fits both fits neither.

“

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

A third option: Build-With

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.⁹ 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.⁹

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.

Source: Kana, *The Build Is Not Your Job*.⁹

05 • BUILD-WITH, CONTINUED

Who holds the maintenance burden three years from now?

The difference between Build-With and the alternatives 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.⁹ Somebody has to update it when the model changes, the data evolves, the regulation shifts, or the platform updates. On a restaurant marketing team, that somebody doesn't exist, and the survey data already told you why: skills are your top-ranked barrier.² We prefer that the technical burden, what we call the technical tax that every enterprise AI deployment carries, belongs to us rather than to you.⁹

“

**Every agent you build in a studio is a divergent artifact
you now own.”**

YOUR CONTROL IS REAL, BUT BOUNDED

YOU DO

Tune outcomes**Set policy****Instrument results**

YOU DON'T

Extend**Fork****Maintain**

Sources: Kana, *The Build Is Not Your Job*; Kana, *The Agentic Divide*.

06 • IN PRACTICE

What agentic marketing means when applied to QSR

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

USE CASE 01 • AUDIENCE BUILDER

Build the segment, not just the demographic

Start here, because it's the fastest proof and the lowest risk. [Kana Audience Builder](#) builds sized, suppressed segments straight from your CRM and loyalty data in minutes rather than through a two-week ticket to the data team. Suppression matters more than sizing in QSR: the guest you shouldn't send the offer to is usually the one who was already coming in. When a segment takes two weeks, the moment that justified it has passed, so the campaign either launches late or launches against a guess.

USE CASE 02 • CAMPAIGN ORCHESTRATOR

Spend that follows the market, not the national average

[Kana Campaign Orchestrator](#) shifts budget toward whatever channel is converting right now, with your team approving every move. High-variance markets stop being stuck on the national mean, and channels that used to run on autopilot finally get a real plan. The weekly comp readout stops being a record of decisions made three weeks too late. This is also the use case that changes the franchisee conversation, because it lets you start from real store-level numbers instead of a system average that nobody's restaurant actually experienced.

USE CASE 03 • PERSONALIZATION

Offer logic built per member, not per program

Given what the retention data says about discount-driven traffic, this is where the margin is. [Kana Personalization](#) builds individualized offer logic per member, so each offer is aimed at adding a visit that wasn't already coming, instead of sending the same buy-ten-get-one to a weekly regular and a lapsed guest alike. We put it third deliberately. It's the most valuable and the most dependent on the two connections above being stable first.

06 • IN PRACTICE, CONTINUED

The order we'd recommend

01

FASTEST PROOF

Audience Builder

Sized, suppressed segments in minutes, not a two-week ticket.

02

CHANGES THE
FRANCHISEE
CONVERSATION

Campaign Orchestrator

Budget follows the markets that are converting, with your OK on every move.

03

WHERE THE MARGIN IS
Personalization

Offer logic per member, aimed at visits that weren't already coming.

On top of your stack, not instead of it

None of this requires data migration. We connect by API and sit on top of Toast, your warehouse, loyalty platform, ordering system, and ad platforms, with no rip-and-replace and no new governance gaps. 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.

KANA WORKS ON TOP OF THE SYSTEMS YOU ALREADY RUN

Audience Builder

Campaign Orchestrator

Personalization

Connected by API • no migration • no rip-and-replace

POS (Toast)

Ordering

Loyalty

Warehouse

Ad platforms

Our team scopes the connection list with you and sets timing per engagement rather than quoting you a number that ignores how your systems are actually wired.

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.⁹ In return, five things are asked of you.

WHAT WE ASK OF YOU

- 01 **Read access to the relevant systems.** If not granted efficiently, it's the single most common source of delay in onboarding.
- 02 A clear account of how audiences are defined, campaigns approved, and co-op and DMA budgets allocated
- 03 Specific and measurable success criteria
- 04 The governance and compliance constraints that apply, including anything your franchise agreements impose on store-level data
- 05 An executive sponsor plus a working lead who will genuinely engage⁹

WHAT YOU RECEIVE

- ✓ A production-grade deployed application, not a prototype or a sandbox
- ✓ A knowledge graph that captures your institutional knowledge and compounds with use
- ✓ Governance training sufficient to own the outcomes
- ✓ Us as the ongoing owner of maintenance as models improve, platforms change, and regulations shift⁹

Restaurant marketing teams have always exercised judgment and taste (pardon the pun). What they have lacked is a marketing operating system worthy of putting that taste to work.⁹

The translation layer between a field marketer'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're ready to help you act on what's happening now.⁹

08 • CONCLUSION

One system, tuned to every market you serve

QSR 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 POS, ordering platform, loyalty program, warehouse, and ad platforms you already run, with no migration, and we run the full lifecycle as a single loop rather than a diagram: building the segment while the moment is still there, checking the kitchen before the offer sends, moving budget toward the markets that are actually converting, and building offer logic per member instead of per program.

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 restaurant organization facing falling traffic, a discounting strategy the retention data has already judged, and franchisees who want proof the spend moved their stores, 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 30-minute demo →](#)

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 ANA, *How Are QSRs and Restaurants Utilizing AI in Their Marketing and Operations?*, 29 July 2026, citing the Qu State of Digital report (survey of approximately 170 limited-service brands) and reported results from Yum Brands and Taco Bell.
- 4 Facteus, *QSR Value Wars Drove Traffic. Most Chains Lost Loyalty*, 8 June 2026. Analysis of 18 major QSR brands using Facteus transaction data, monthly retention rates from May 2024 through April 2026.
- 5 QSR Magazine, *The State of Restaurant Traffic as 2026 Gets Underway*, 16 January 2026, citing Revenue Management Solutions Q4 2025 data.
- 6 eMarketer, *Restaurant chains need sharper value, not just more promotions*, 5 August 2026, citing Placer.ai.
- 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, *Evaluation Guide: Is Your AI Enterprise Grade?*
- 9 Kana, *The Build Is Not Your Job*, June 2026.

QSR capability descriptions, integration details, and franchise contribution ranges are drawn from Kana's published QSR materials.

ABOUT KANA

Kana is a suite of agentic marketing applications built around your business goals and the systems you already run. It brings scattered customer and business data together, then puts digital helpers to work on the busywork, so marketing teams can make smarter decisions and act on them faster. Nothing goes live without your OK, and your data stays in your systems, under your control.

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