

Moving at the Speed of Your Pricing

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

CONTENTS**INSIDE THIS PAPER**

—	Executive summary	3
01	Travel'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 travel and hospitality	11
05	The build is not a travel marketer's job	14
06	What agentic marketing means when applied to travel and hospitality	16
07	What a Build-With engagement requires	18
08	Conclusion	19
—	Endnotes	20

EXECUTIVE SUMMARY

LATE DECISIONS ARE THE MOST EXPENSIVE KIND

Travel and hospitality has a timing problem that no amount of incremental tooling has been able to resolve.

A competitor drops rates overnight, the rate shopper alerts revenue at 6am, and marketing hears about it on Thursday. A guest checks out on Sunday and the post-stay email fires on a timer that was set months ago, whether or not that guest is anywhere near ready to book again. A top-tier member who has stayed with you all year receives the same offer as someone who subscribed yesterday.

Each of those gaps is a revenue decision made late, and in an industry where pricing moves by the hour and an unsold room night can never be sold again, late decisions are the most expensive kind of decision that travel 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. Travel and hospitality is where that architecture matters most, because demand is perishable, because the guest relationship is increasingly contested by intermediaries, and because the discovery layer that feeds travel 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 travel and hospitality organizations really are to adopt agentic AI
- 03** What enterprise-grade has to mean when the data in question is your guest profiles, your loyalty program, and your rate strategy
- 04** How agentic marketing applies to the specific decisions travel and hospitality marketing teams make every week

This paper was written for travel and hospitality 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 travel marketing. Let's dive in.

01 · THE TIMING PROBLEM

TRAVEL'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 travel) 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 hotel groups and travel brands have accumulated at least a property management system, a central reservation system, a revenue management system, a booking engine, a loyalty platform, a CRM, an ESP, a rate shopper, metasearch and paid media accounts, and an analytics layer that reconciles none of them very 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
judgment

80%+

of hotel commercial
teams still spend one
to two days a week
on manual reportingSources: Kana, Understanding AI Agents for Marketing;¹ The State of Distribution 2026.⁵

The critical issue is that 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 marketer became that connective tissue. 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 Hospitality's own benchmark data says the same thing in its own terms: more than 80 percent of hotel commercial teams still spend one to two days a week producing and analyzing reports manually.⁵


As stated above, that is an architecture problem, not a tooling problem, and it shows up in travel 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 email sent on a fixed schedule to a guest who already rebooked somewhere else. It appears as a campaign still promoting yesterday's rate after a competitor has moved. It appears as loyalty spend that cannot be tied to incremental revenue when the CFO asks for proof.

01 · THE TIMING PROBLEM, CONTINUED

TWO FORCES MAKING THE LAG LESS SURVIVABLE

The guest relationship is increasingly owned by someone else. Across more than 58,000 properties, online travel agencies (OTA) now generate close to twice the bookings of hotel-owned digital channels,⁵ with OTAs at 34 percent of bookings in 2026 against 18 percent for hotel websites and apps.⁶ Among independent hotels, OTA share rose to 63.4 percent in 2025, and between 2019 and 2025 the cost of acquisition grew 25 percent while revenue per available room grew 19 percent.⁷ The channel matters long after the first stay: in one analysis of 2.2 million guests at European hotel groups, guests who booked direct returned at 5.1 percent against 1.3 percent for OTA-acquired guests, and OTAs supplied 65 percent of first stays but only 32 percent of return stays.⁸

WHO OWNS THE BOOKING, AND WHO COMES BACK



5.1%

return rate for guests who booked direct

1.3%

return rate for OTA-acquired guests

Sources: PhocusWire, citing The State of Distribution 2026 and H2c;⁶ eHotelier, citing Bookboost (2.2 million guests).⁸

Meanwhile, the loyalty programs meant to protect that relationship are getting more expensive. In 2024 U.S. hotel loyalty program fees grew 3.9 percent, outpacing growth in both rooms revenue (2.7 percent) and rooms occupied (1.3 percent), even as loyalty members' contribution to occupancy rose to 52.8 percent and average room nights per member fell to 1.0.⁹

LOYALTY COSTS ARE OUTRUNNING THE REVENUE THEY PROTECT, U.S. HOTELS, 2024

+3.9%

growth in loyalty program fees

+2.7%

growth in rooms revenue

+1.3%

growth in rooms occupied

Source: CBRE Hotels Research.⁹

01 · THE TIMING PROBLEM, CONTINUED

TRAVEL DISCOVERY IS BEING REBUILT AROUND AI

Travel discovery is being rebuilt around AI faster than most travel measurement has adjusted. In early 2026, 56 percent of U.S. travelers had used AI for planning, booking, or in-destination help on at least one trip in the past 12 months, up from 33 percent in the first half of 2025.¹⁰ Traffic from AI sources to U.S. travel sites grew 194 percent year over year in May 2026 and is up 2,215 percent since October 2024. Those AI-referred visitors spent 70 percent longer on site, were 21 percent more engaged, and bounced 41 percent less than visitors from other sources, and while they still converted 28 percent below non-AI traffic, that gap has narrowed by nearly 70 percent since late 2024.¹¹

AI-REFERRED VISITORS TO U.S. TRAVEL SITES, MAY 2026

+194%AI-referred traffic,
year over year**+70%**longer time on site
than other sources**-41%**lower bounce rate
than other sources**56%**of U.S. travelers used
AI on at least one trip
in the past yearSources: Search Engine Land, citing Adobe Analytics;¹¹ Phocuswright Research via PhocusWire.¹⁰

AI search already accounts for 4 percent of hotel reservations, yet 55 percent of hotels report no change or only a minor change to their distribution strategy.⁶ And in September 2026, Meta launched an AI agent that can plan and book travel on a user's behalf, with accommodation bookings expected to follow.¹²

A channel that moves this quickly is not something that a plan developed and deployed annually can necessarily track and keep pace with.

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 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

01 The model Supplies reasoning.	02 The prompt Encodes your objectives, brand rules, and constraints.	03 Memory Accumulates business context so the system improves the longer it runs.	04 Tools The systems and APIs the agent can act through, which give it hands and not just a brain.	05 History The complete record of what it decided and why: a learning mechanism and an audit trail. ¹
--	--	---	--	--

02 · DEFINING THE AGENT, CONTINUED

THREE LEVELS OF CAPABILITY, IN TRAVEL TERMS

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

01

START HERE

On-demand agents

Activate when asked and complete a defined task. In travel, that can mean consolidating performance across the booking engine, the PMS, the loyalty file, and paid media 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 an overnight competitor rate drop and handing marketing the alert and the evidence before the morning stand-up, or flagging a cohort of loyalty members whose booking frequency has started to decay.

03

AS TRUST BUILDS

Autonomous agents

Pursue an outcome with minimal direction, designing, deploying, monitoring, and iterating a lapsed-member reactivation 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 Travel teams that invert that order may find that they discover their guest-profile and data problems at the worst possible moment, which is usually the run-up to peak season.

03 · THE READINESS GAP

ADDRESSING THE READINESS GAP

The market has already decided that agentic marketing is happening. In a survey of 400 hospitality IT decision-makers, half of hotels said they are adopting AI now and 82 percent plan to ramp up usage in the next 12 months.⁷ Consumer-facing AI agents are already booking travel.¹² 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

50%

of hotels say they are adopting AI now⁷

82%

of hotels plan to ramp up AI usage in the next 12 months⁷

82%

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

46%

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

Sources: PhocusWire, citing Canary Technologies (n=400); 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 The State of Distribution 2026, drawn from more than 270 hotel brands and 58,000 properties, which found that more than half of hotels now use or are procuring generative AI, yet fewer than one in ten say it has reduced their manual work by more than 30 percent.⁵ Separate research found that 69 percent of hotels still have staff manually entering guest preference data, compared with 19 percent extracting it from guest interactions with AI.⁶ Both numbers cannot be describing the same condition.

READINESS, AS REPORTED VS. AS EXPERIENCED

Leaders who call their data infrastructure ready²




86%

Hotels still entering guest preference data by hand⁶




69%

Hotels extracting guest preferences with AI⁶



19%

Hotels where AI cut manual work by 30%+⁵



<10%

Sources: Kana, The Agentic Divide; The State of Distribution 2026; PhocusWire, citing H2c.

03 · THE READINESS GAP, CONTINUED

A CEILING ON AUTONOMY, AND WHERE IT COMES FROM

The State of Distribution also found that trust, data privacy, and governance gaps are creating a ceiling on AI autonomy: hotels are increasingly willing to use AI for assistance but remain cautious about letting it make decisions or act on its own.⁵ Practitioners describe the root cause plainly. As one hotel group's CRM and loyalty lead put it:

Duplicated guest profiles, inconsistent preferences, and data spread across systems do not disappear when AI is layered on top. They get amplified.

Read together, these findings describe a specific but avoidable failure mode. Travel and hospitality organizations are adopting AI quickly, using it mostly to make existing tasks faster, assessing their own foundations generously, and hesitating to trust it with decisions because the governance to support that trust was never built. That is how a pilot becomes a stalled program.

But it is also why the question a travel CMO should be asking is not which AI features a vendor offers, but **what the vendor is prepared to own**. We'll explore that in more detail in the next section.

“

AI is only as useful as the data and processes behind it.⁶

CRM AND LOYALTY LEAD, HOTEL GROUP

WHAT ENTERPRISE GRADE HAS TO MEAN IN TRAVEL AND HOSPITALITY

Most enterprise AI failures are not model failures. They are the result of missing context, weak governance, shallow evaluation, or fragile integrations.³ In travel and hospitality 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 travel marketing 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 hotel group or travel brand, that means the PMS, CRS, RMS, booking engine, loyalty platform, CRM, ESP, and ad platforms, connected without a migration.

PMS

CRS

RMS

Booking engine

Loyalty platform

CRM

ESP

Ad platforms

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 travel and hospitality, that business is your property portfolio, your rate structure and fences, your loyalty tiers and their economics, your seasonality and event calendar, and your channel mix, none of which a general model has ever seen.

“

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

GOVERNANCE, AUDIT AND EVALUATION

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: travel teams handling guest profiles, loyalty accounts, and payment-adjacent booking 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 member offer 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. Travel's seasonality guarantees that a model validated in shoulder season will meet peak-season demand, event-driven spikes, and disruption it has never seen, and it must be built to anticipate that.

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 travel 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 weekly revenue meeting.³



Travel teams handling guest profiles, loyalty accounts, and payment-adjacent booking data should treat pre-action enforcement as a threshold requirement rather than a feature.

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. That is exactly the posture hotel commercial teams say they want today: assistance first, autonomy earned.⁵ And recommendations have to activate: a recommendation that stops at a dashboard has not been delivered.³

“

**Assistance first,
autonomy earned.**

The six tests of enterprise grade, at a glance

01**Integration that writes
as well as reads**

Data in and data out across the PMS, CRS, RMS, booking engine, loyalty, CRM, 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, through every season.

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 TRAVEL MARKETER'S JOB

Travel and hospitality marketing leaders are usually offered two options for procuring new software.

- **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 travel brands, 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 Hotel commercial teams appear to have reached the same conclusion: their leading technology priorities are improving existing systems and reducing integration and vendor complexity, rather than adding new platforms.⁵

“

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 travel brand's workflow, configured to its data, governed by its policies, and tuned to its tolerances, at the speed and cost of buying, with the durability of a vendor relationship.⁴

BUILD-WITH IS NOT**Low-code**

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

BUILD-WITH IS NOT**Consulting**

Consulting produces one-off systems that rot.

BUILD-WITH IS 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

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 hotel or travel marketing team, that someone simply does not exist.

“

We believe the technical burden, what Kana calls the technical tax that every enterprise AI deployment carries, belongs to Kana, not to the customer.⁴

Shared infrastructure no travel brand should solve alone

Underneath the applications sits shared infrastructure that no individual travel brand 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

Governance and guardrails

LLM routing and cost optimization

Attribute- and role-based access control

Nondeterminism management

Knowledge graph compilation

Evaluation and observability

Data connectors in and out

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

— 06 · IN PRACTICE

WHAT AGENTIC MARKETING MEANS WHEN APPLIED TO TRAVEL AND HOSPITALITY

So how will agentic marketing be successfully applied in a travel and hospitality context? What are the best low-hanging fruit use cases where the impact of agentic marketing can be most immediately felt?

01 — USE CASE · PERSONALIZATION

Stop sending first-timer offers to your best members

Personalization stuck in batch segments means high-value members get the same email as brand-new subscribers. That is a costly habit when loyalty program fees are growing faster than rooms revenue,⁹ and when batch segments cannot prove the incremental revenue the CFO wants to see. Kana's Personalization application reads recent member activity, including browsing, app opens, stays, and tier status, and recommends each guest's next offer for the team to approve. Top-tier members stop receiving first-timer emails, and loyalty spend starts producing revenue that can be traced. Every offer still follows the program's approval rules and can be subject to human reviews and approvals.

02 — USE CASE · CATEGORY INTELLIGENCE

Hear about the rate move before sales does

Rate changes still require manual monitoring for most marketing teams. A competitor drops rates and marketing hears about it from sales, days late, often after a campaign promoting the old price has already gone out. Kana Category Intelligence watches competitive pricing daily and alerts marketing the moment something shifts, with the evidence attached, rather than relaying a secondhand report. The revenue team's rate shopper watches prices for revenue management; Category Intelligence watches them for marketing, so messaging and offers can respond the same day.

03 — USE CASE · AUDIENCE BUILDER

Build the list worth re-engaging

Win-back and anniversary lists are usually static exports that were already out of date when they were pulled. Kana Audience Builder builds precise, suppressed guest segments directly from CRM, loyalty, and booking data, including lapsed members, recent cancellations, and high-value repeat bookers, so outreach reflects who is actually eligible today. That precision matters in a market where nearly 22 percent of OTA bookings at independent hotels are cancelled, roughly double the 10.6 percent rate across independent hotel bookings,⁷ and where every OTA-acquired guest is a relationship the property has to earn back before the second stay.⁸

FROM THE BOOKING WINDOW TO THE READOUT

04 — USE CASE · AUDIENCE BUILDER + PERSONALIZATION Meet the guest's booking window, not the calendar

Outreach in most hotel stacks runs on a calendar rather than on the guest's window. Post-stay and win-back sequences fire on fixed timers regardless of what the guest does next, and coordinating email, app, and paid outreach means stitching three tools together by hand. The team often knows which moments convert best but cannot trigger outreach at that moment instead of the scheduled one. With Kana Audience Builder deciding who to reach and Kana Personalization deciding what to offer and when, outreach is timed to when a guest is ready to book again, not to when the automation was configured to send.

05 — USE CASE · CAMPAIGN ORCHESTRATOR Catch spend drift before the readout

Campaign performance in travel is still largely reviewed in arrears, which means a channel drifting from the brief can burn budget for weeks before anyone sees it. Kana Campaign Orchestrator watches live spend across channels, flags any channel that drifts from the brief, and proposes a reallocation for approval before the monthly readout catches it two weeks too late.

On top of the stack, not instead of it

None of the above requires a rip-and-replace or a data migration. Kana connects by API and sits on top of the PMS, RMS, CRS, loyalty platform, and booking engine a travel brand already runs, including Opera, Duetto, IDEaS, and Salesforce. In some ways you could say we're actually here to help you get more return on your prior martech and adtech investments.

Kana's team handles setup, and timelines depend on how many systems are connected. Most travel and hospitality teams see value within 30 days by connecting their top three to five data sources, starting with Personalization on loyalty data and Category Intelligence on the rate topics they want watched, then layering in Audience Builder and Campaign Orchestrator once those connections are stable and insights and decisions are being delivered routinely and reliably.

FROM CONNECTION TO VALUE

3–5

top data sources connected to start

30 days

for most travel and hospitality teams to see value

WHAT A BUILD-WITH AGENTIC MARKETING ENGAGEMENT REQUIRES

Build-With is a commitment rather than a capability, and it's worth being explicit about [what it requires from a travel and hospitality 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 TRAVEL BRAND

- 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, offers and campaigns approved, and loyalty and win-back 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 TRAVEL BRAND 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

Travel and hospitality marketing has always had judgment. What it has lacked is a system worthy of it.⁴ The translation layer between a marketer's judgment and the systems that execute on it has been staffed by people, governed by ticket queues, and measured in days, while pricing moves by the hour. Collapsing that layer is the point. Kana will help you act on what is happening now.⁴

08 · CONCLUSION

ONE SYSTEM, MOVING AT THE SPEED OF YOUR PRICING

Travel and hospitality 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 PMS, RMS, CRS, loyalty platform, and booking engine you already run, with no rip-and-replace and no migration, and run the guest lifecycle as a single loop rather than a diagram: building the segments guests actually belong in, matching each member's offer to the tier and activity they have earned, timing outreach to the booking window instead of the calendar, and putting the rate move in front of marketing the same day it happens.

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 Kana rather than landing on a marketing team that was never staffed to pay it.

For a travel or hospitality organization facing intermediaries that own a growing share of the guest relationship, loyalty costs that are rising faster than room revenue, and a discovery layer being rebuilt around AI agents, 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 a 30-minute call to discuss what this could 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 NYU School of Professional Studies, RateGain and HEDNA, [More Than 50% of Hotels Use AI, but Under 10% See Real Impact, Finds State of Distribution 2026 Report](#), 15 September 2026. Based on more than 270 hotel brands and 58,000+ properties across 53 countries; survey period December 2024 to November 2025.
- 6 PhocusWire, [Hotels still stuck on manual work despite increased AI adoption](#), 1 October 2026, citing The State of Distribution 2026 (NYU SPS, RateGain, HEDNA) and H2c research.
- 7 PhocusWire, [Independent hotel reliance on OTAs increased in 2025](#), 25 March 2026, citing the Cloudbeds State of Independent Hotels study (90 million bookings) and Canary Technologies research (400 hospitality IT decision-makers).
- 8 eHotelier, [Direct booking guests return at 4x the rate of OTA guests: what the data shows](#), 9 September 2026. Bookboost analysis of 2.2 million guests at European hotel groups, 2024 to 2026 (vendor-published data).
- 9 CBRE Hotels Research, [For Hotel Owners Loyalty Program Benefits and Costs Continue to Rise](#), 29 October 2025. Analysis of 4,187 U.S. hotels, 2023 to 2024, and public filings of Choice, Hilton, Hyatt, Marriott and Wyndham.
- 10 Phocuswright Research via PhocusWire, [The fastest shift in travel behavior just became the default](#), 24 March 2026. Summary of Phocuswright's The AI Surge: Travel's Fastest Behavioral Shift in a Decade, a study of U.S. travelers.
- 11 Search Engine Land, [AI referrals to travel sites surge 194% as engagement rises: Adobe](#), 17 June 2026, citing Adobe Analytics data for May 2026 based on more than 8 million visits to U.S. travel sites.
- 12 PhocusWire, [Meta launches AI agent with travel booking capabilities](#), 9 September 2026.

Travel and hospitality capability descriptions and integration details are drawn from Kana's published travel and hospitality materials. [kana.ai](#)