

What AI Knows About the Traveler: How LLM Personalization Is Reshaping Travel Discovery, Research, and Booking

A research report for hospitality executives, destination organizations, and travel marketers

David Metze, VP AI Innovation - July 2026


Executive Summary

The most consequential shift in travel marketing since the rise of Google is now underway, and it has as much to do with what AI assistants know about the world as with what they know about the individual traveler.

Between 2024 and 2026, every major consumer AI assistant — ChatGPT, Gemini, Claude, Copilot, Perplexity, and Grok — moved from stateless, session-based chat to **persistent, memory-enabled personalization**. Several now connect directly to users' email, calendars, photos, search history, and purchase-adjacent data. The result is that two travelers typing the identical prompt — "Where should we stay in Maui in February?" — can receive materially different answers, shaped by inferred budget, family composition, loyalty affiliations, dietary needs, past trips, and even the tone of prior conversations.

For hotels, resorts, and destinations, this changes the physics of discovery. Traditional search presented a ranked list of links to a largely anonymous user; the traveler did the filtering. AI assistants now do the filtering *before* the traveler sees anything, using a private, invisible profile the marketer cannot observe, buy, or directly target. Visibility is no longer about ranking on a results page. It is about being the answer for a specific person — which means being legible, trustworthy, and richly described enough that an AI can confidently match a property to a known set of preferences.

Adoption is no longer speculative. Depending on the survey, between roughly 30% and 65%+ of travelers now use AI tools somewhere in trip planning, with penetration above 60% among the affluent, Millennials and Gen Z, and openness above 90% in markets like India and China. AI-referred traffic remains a small share of hotel website visits (roughly 1% in mid-2026 per Lighthouse data) but is growing extremely fast — up more than 50% within weeks of ChatGPT expanding outbound links — and it converts at meaningfully higher rates than generic organic search because AI users arrive pre-qualified and pre-decided. Meanwhile, OpenAI's launch (and partial retreat) of in-chat commerce, the arrival of Expedia, Booking.com, Tripadvisor, and Accor apps inside ChatGPT, and Google's announced agentic booking partnerships with Marriott, IHG, Choice, and Wyndham all point in one direction: the AI conversation is becoming the top of the travel funnel, and increasingly the middle of it too.



Five conclusions for hospitality leaders:

1. **Personalization moves the point of persuasion upstream and out of sight.** The AI's private model of the user — not the ad auction, not the SERP — increasingly determines which three to five properties a traveler ever considers. You cannot target this profile; you can only be maximally matchable against it.
2. **Signal-rich, machine-readable content is the new media buy.** AI assistants match user attributes (family status, accessibility needs, dietary preferences, loyalty membership, budget tier) to property attributes. Properties that publish specific, structured, verifiable detail about who they serve win disproportionate share of personalized recommendations.
3. **First-party data and loyalty are appreciating assets.** As assistants become agents that transact, they will pull traveler profiles, loyalty status, and preferences from the systems the traveler authorizes. Hotels that own the guest relationship and expose it via clean APIs will be the "authoritative source" agents rely on; those that don't will be intermediated again.
4. **The evidence is early and uneven.** Full autonomous booking remains rare (only ~2–5% of travelers say they would let AI book without oversight, per Skift/McKinsey research), OpenAI walked back in-chat Instant Checkout for travel in early 2026, and measurement of AI influence is immature. Leaders should invest ahead of the curve without over-rotating on any single platform.
5. **The window is now.** AI referral traffic is currently commission-free, AI answer sets are still forming, and early movers are disproportionately cited. The cost of experimentation is low; the cost of invisibility compounds.

The remainder of this report explains how each major assistant learns about users, how that knowledge reshapes recommendations across every travel category and journey stage, which personalization signals matter most, and what independent hotels, luxury brands, DMOs, agencies, revenue managers, and technology providers should do about it between now and 2030.

1. How LLMs Learn About Users

1.1 The layered anatomy of AI "knowledge" about a person

It is tempting to imagine an AI assistant's knowledge of a user as a single dossier. In practice, personalization is assembled fresh at every message from several distinct layers, each with different persistence, visibility, and privacy characteristics. Understanding these layers is essential, because they determine *what a travel recommendation can be conditioned on* — and what marketers can and cannot influence.

Layer 1: The context window (temporary, session-scoped). Every LLM response is generated from a "context window" — the working memory containing the current conversation, any uploaded files, and system instructions. Anything the user says in-session ("we're a family of



five," "my wife uses a wheelchair," "we loved the Amalfi Coast") immediately shapes recommendations, but this raw context evaporates when the session ends unless a memory system captures it. Independent reverse-engineering of ChatGPT's architecture describes a structured context assembled per message: system and developer instructions, session metadata (device, timezone, approximate location), stored user memories, summaries of recent conversations, and finally the live chat. Session metadata is ephemeral; it explains why assistants can reference "this morning" correctly or adapt to mobile without "remembering" anything.

Layer 2: Explicit, user-stated persistent memory. All major assistants now support deliberate memory: "Remember that I'm vegetarian," "Remember I fly out of Denver," "Remember my kids are 6 and 9." These saved memories are injected into future conversations and typically viewable and deletable in settings (e.g., ChatGPT's Settings → Personalization → Manage Memories; Gemini's saved-info; Claude's editable memory summary). Custom instructions and profile fields (name, role, preferences) belong to this layer too. This is the highest-signal, most reliable personalization input — and, notably, one users can be *coached* to create (more on this in Section 6).

Layer 3: Derived / inferred memory from conversation history. Since 2025, the leading assistants also synthesize insights from past chats without being asked. ChatGPT's "chat history" reference (expanded April 2025) draws on all prior conversations to build a longer-term understanding of Plus/Pro users. Gemini's Memory learns "about you and your world" from prior chats. Claude generates a rolling synthesis of chat history, updated roughly every 24 hours, with separate memory spaces per project. Perplexity's memory (launched November 2025, upgraded through 2026) automatically synthesizes preferences, favorite brands, dietary needs, and recurring topics, retrieving them from a private memory store rather than treating history as undifferentiated data. Microsoft's Copilot Memory reached general availability across Microsoft 365 in early 2026, storing preferences and working style inside the user's Exchange mailbox with enterprise-grade controls. Grok enabled persistent memory by default outside the EU/UK in 2025 and added reusable "Skills" in 2026. This layer is powerful but interpretive: memories are LLM-generated summaries, which can be stale or wrong — a stored "lives in Manchester" can silently shape recommendations months after a move.

Layer 4: Connected applications and ecosystem data. This is where personalization depth diverges sharply by vendor. Google's Personal Intelligence (rolling out from late 2025/2026, excluding the EEA, UK, and Switzerland) lets Gemini draw on Gmail, Calendar, Drive, Google Photos, Search/Maps/Shopping history, and YouTube history — explicitly marketed with travel examples like tailored itineraries and pulling photos "from your last vacation." A Gemini user's flight confirmations, hotel folios, restaurant reservations, and past searches constitute an extraordinarily rich implicit travel history. Microsoft Copilot similarly reads the user's work graph (Outlook mail, calendar, OneDrive documents), which is decisive for business-travel context. ChatGPT supports connectors and in-chat apps (Expedia, Booking.com, Tripadvisor, and others) that inject live inventory and, with authorization, account context. Claude connects to tools and files via integrations but does not mine a consumer data ecosystem. Grok is tied to X, giving it real-time social context rather than transactional history.

Layer 5: Location, account, and behavioral signals. Assistants generally know coarse location (from IP or device), account identity, language, subscription tier, and interaction patterns (what the user asks about repeatedly, how they respond to suggestions, which recommendations they

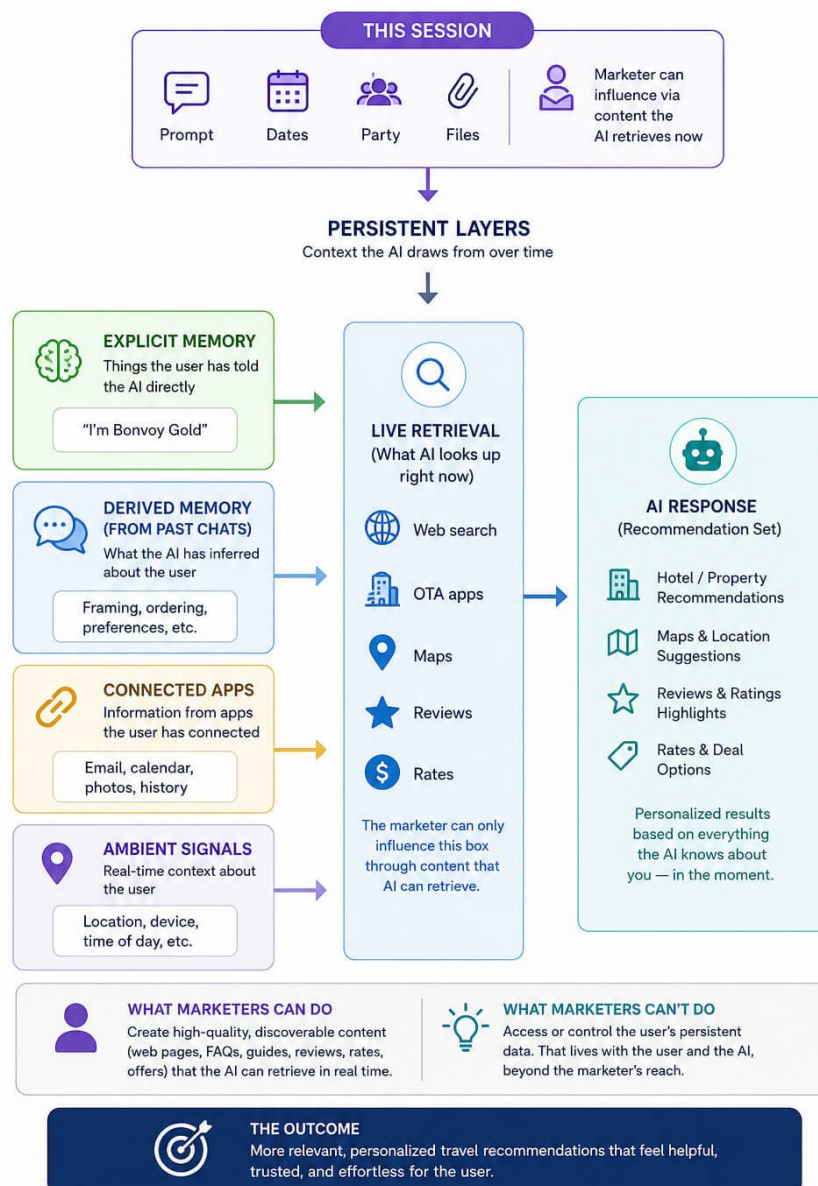


pursue). Purchase history per se is not natively visible to most assistants — but it becomes visible the moment a user connects Gmail (receipts, confirmations) or an OTA account, and agentic checkout systems are beginning to formalize this.

1.2 Temporary context versus persistent personalization

The critical distinction for travel marketers is between **conversational context** (what the user tells the AI during this planning session — dates, destination, party size) and **persistent personalization** (what the AI already believed about the user before the session began). Conversational context makes any AI good at *this* trip. Persistent personalization changes *which options are ever surfaced*: the default budget band assumed, the brands mentioned first, whether "family-friendly" or "adults-only" is presumed, whether the assistant volunteers accessibility information unprompted.

How User Knowledge Flows Into an AI Travel Recommendation



1.3 What is and is not retained — and privacy limits

Retention is configurable and vendor-specific, but common patterns hold across platforms. Users can typically view, edit, and delete saved memories; disable memory entirely; and use private modes (ChatGPT Temporary Chat, Claude Incognito, Perplexity/Grok incognito, Gemini per-chat Personal Intelligence toggle) that neither read nor write memory. Deleting a chat does not necessarily delete memories derived from it (true in both ChatGPT and Gemini documentation), a nuance with real privacy implications. Enterprise and business tiers generally exclude memory content from model training and add admin-level kill switches.

Regulation meaningfully shapes availability. Gemini's Personal Intelligence launched excluding the EEA, UK, Switzerland, and Nigeria; ChatGPT's chat-history referencing arrived later and more cautiously in Europe; Grok's default-on memory excludes the EU and UK. Commentators have linked these gaps to GDPR requirements around transparency, purpose limitation, and data portability — the derived-profile layer is exactly the kind of opaque inference GDPR scrutinizes. The practical consequence: **the depth of travel personalization a guest experiences already varies by geography**, and European travelers will, for now, encounter shallower default personalization than American or Asian ones.

Two further limits deserve executive attention. First, memories are *summaries*, not recordings — assistants store interpretive notes ("prefers boutique hotels, budget-conscious, two young children") rather than transcripts, which conserves privacy but introduces error. Second, none of the consumer assistants sells or exposes these profiles to advertisers; several (including Anthropic) have publicly committed to ad-free assistant experiences. There is currently **no mechanism for a hotel to buy access to, or targeting against, an AI's user profile**. Influence must come through content, data, reputation, and integrations.

2. How User Knowledge Changes Travel Recommendations

2.1 The mechanics of personalized recommendation

When an assistant answers a travel query, personalization operates at four points: **interpretation** (what the vague word "nice" means for this user's budget), **candidate generation** (which properties are retrieved or recalled at all), **ranking and framing** (which of the candidates lead, and which attributes are emphasized), and **omission** (what is silently filtered out — often the most consequential and least visible step). ***A property can be perfectly optimized for generic AI search and still never appear for a given traveler because the assistant "knows" that traveler avoids high-rises, or resorts, or that brand.***

2.2 One prompt, three travelers: a worked illustration

Consider the identical prompt sent to a memory-enabled assistant by three different users:

"We're thinking about a week in Maui in February. Where should we stay?"



	User A	User B	User C
What the AI knows	Prior chats about kids' swim lessons and a 529 plan; saved memory "two children, 6 and 9"; past trip questions about Disney World; Gmail shows Southwest and Marriott confirmations; asked last month how to use Bonvoy points	Frequent chats about restaurant tasting menus and wine; calendar shows an anniversary in February; past stays referenced at Aman and Rosewood properties; no children ever mentioned; high-spend signals	Chats about trail running and surf forecasts; previously asked about hostels and vanlife; budget questions recur; vegetarian noted in saved memory; travels solo
Likely interpretation	Family resort trip, points-relevant, kid programming matters	Romantic luxury escape, food-led, privacy-oriented	Active budget trip, location over amenities
Likely recommendation set	Family-oriented resorts in Kā'anapali/Wailea with kids' clubs, pools, connecting rooms; Bonvoy-affiliated options surfaced first; note on redeeming points	Adults-oriented luxury in Wailea; suite/villa options; chef's-table and sunset-sail suggestions; no mention of kids' clubs	Condos/guesthouses near Pā'ia; food-truck and hike recommendations; car-rental and shoulder-pricing tips
What gets omitted	Adults-only boutiques; hostels	Family mega-resorts; value chains	Full-service luxury resorts

No user asked to be segmented. The segmentation happened invisibly, before the first option was named. This is the core structural difference from a Google search, where all three users would have seen substantially the same results page and self-filtered.

2.3 Category-by-category effects

Hotels and resorts. Persistent knowledge shifts the assumed tier (budget/midscale/upper-upscale/luxury), the brand set (loyalty memberships pull affiliated properties forward), and the attribute emphasis (kids' club vs. spa vs. meeting space). Because AI answers typically name only three to five properties, personalization concentrates demand: the "consideration set" is now assembled by the machine.

Vacation rentals. Family composition, pet ownership, length of stay, and kitchen/dietary signals push assistants toward rentals over hotels for some users — or vice versa for users who have complained about rental check-in friction. A remembered "we travel with a golden retriever" quietly reorders the entire lodging market for that household.



Cruises. Prior cruise mentions (line, ship class, cabin type), motion-sickness or accessibility notes, and multigenerational-travel context shape line and itinerary suggestions. Cruise demand is rising (30% of travelers in IMG's 2026 survey plan a cruise, up 11% year over year), and cruise selection is complex enough that AI synthesis is unusually valuable — and unusually decisive.

Airlines and airports. Preferred home airport, alliance status, seat and cabin preferences, and tolerance for connections are among the stickiest remembered facts. Assistants with calendar/email access (Gemini, Copilot) can infer route history directly from confirmations.

Destinations. Climate preferences, risk tolerance, visited-before history, passport/visa context, and values signals (sustainability, crowd-aversion) reshape which destinations are proposed at the inspiration stage — the stage DMOs care most about, and the one where AI influence is strongest today (survey after survey shows discovery/ideation as the top AI use case).

Restaurants, tours, activities, attractions. Dietary preferences and allergies (among the most commonly saved memories), mobility constraints, kid ages, and stated interests (art, golf, diving, wine) personalize the in-destination layer. Perplexity explicitly markets dietary-need memory; Gemini can cross-reference Maps and Search history.

Luxury travel. Personalization amplifies luxury discovery in two ways: assistants infer spend capacity from accumulated context rather than requiring the user to declare it, and they emphasize the *specific* luxury dimensions the user cares about (privacy vs. status vs. gastronomy vs. wellness). Undifferentiated "5-star" positioning performs poorly here; sharply articulated identity performs well.

Family travel. Kid ages are gold-standard signals — programming for a 4-year-old and a 15-year-old barely overlaps. Assistants that remember children's ages produce dramatically better (and stickier) family recommendations, and OpenAI's own marketing examples for memory feature exactly this pattern (remembering a toddler's interests).

Business travel. Copilot and Gemini can see the meeting on the calendar, the client's address in the email thread, and corporate travel-policy documents. Recommendations become constraint-solving: within policy, near the meeting, on the preferred alliance, with the loyalty number attached. This is the segment where connected-ecosystem assistants have the largest structural advantage.

Wellness travel. Fitness routines, injuries, sleep issues, and dietary regimes discussed in past chats (Perplexity's running-shoe example generalizes directly) let assistants match travelers to specific programming — a hot-springs recovery property vs. a bootcamp vs. a medical spa.

Golf travel. Handicap conversations, home-course mentions, and equipment questions signal skill level and seriousness; assistants can match course difficulty, walkability, caddie culture, and stay-and-play packaging accordingly.

Meetings and events. For planners, assistants remember organization size, past venues, F&B constraints, AV requirements, and budget bands across RFP cycles — turning venue search into an increasingly automatable procurement task, with implications for hotel group-sales teams (respond to machine-structured RFPs, publish capacities and floor plans in machine-readable form).



3. Personalization Signals, Ranked

Not all user attributes influence recommendations equally. Based on how the platforms document their memory systems, how recommendation quality responds to each attribute, and how frequently each attribute is captured in practice, the following ranking reflects our assessment of **impact on travel recommendation quality and outcome**. (This is informed analysis; no platform publishes signal weights.)

Rank	Signal	Why it matters	Capture likelihood today
1	Stated preferences & explicit memories	Direct, unambiguous, always injected into context; overrides inference	High — all platforms
2	Budget / luxury-vs-value orientation	Gates the entire candidate set; wrong tier = useless answer	High — inferred quickly from language and history
3	Family status & children's ages	Reorders lodging type, room config, destination, activities	High once mentioned; sticky
4	Travel history (destinations, properties, trip styles)	Best predictor of future taste; enables "more like X, but..."	Medium — high with connected email/photos
5	Geographic location & preferred airports	Determines feasibility, pricing, and nonstop options	High — ambient signal plus memory
6	Accessibility needs	Hard constraint; assistants treat as non-negotiable filter	Medium — undercaptured; huge quality gain when present
7	Loyalty memberships & preferred brands	Pulls affiliated properties forward; decisive at booking/agentic stage	Medium — rising as agents integrate accounts
8	Dietary preferences & allergies	Shapes restaurants, resorts (AI/all-inclusive), cruises	High — among most commonly saved memories
9	Travel style (pace, boutique vs. resort, adventure vs. rest)	Governs framing and ordering within the tier	Medium — accretes across conversations
10	Seasonality & climate preferences	Redirects destination and timing suggestions	Medium
11	Previous conversations (topic-level)	Ambient context that colors interpretation of vague prompts	High on memory-enabled platforms
12	Pets	Binary filter with outsized effect on lodging set	Low-medium — but decisive when present
13	Risk tolerance & health context	Adjusts destination safety framing, insurance suggestions	Low-medium
14	Income assumptions	Usually mediated through budget/behavior rather than used directly	Inferred, rarely explicit
15	Implicit behavioral patterns (session behavior, engagement)	Fine-tunes tone and depth more than the option set	Platform-internal, opaque

Three strategic observations follow. First, **the top signals are ones travelers volunteer** — which means brands can legitimately encourage capture ("tell your AI assistant you're a Vail



Resorts Epic Pass holder" is a plausible future CRM message). Second, **constraint signals (accessibility, pets, dietary) punch above their frequency**: they act as hard filters, so properties that publish specific, verifiable information on these dimensions capture near-monopoly share of the filtered answer set. Third, **loyalty is about to jump up this table**: as assistants become booking agents that authenticate into member accounts (Google's announced agentic booking names Marriott, IHG, Choice, and Wyndham as partners), loyalty status becomes a live transactional signal rather than a remembered fact.

4. Platform Comparison: Personalization Capabilities of the Major LLMs

4.1 Comparison table (as of mid-2026)

Capability	ChatGPT (OpenAI)	Gemini (Google)	Claude (Anthropic)	Perplexity	Grok (xAI)	Copilot (Microsoft)
Explicit saved memories	Yes; manageable in settings; memory summary view	Yes ("remember..."); saved info + instructions	Yes; editable memory summary; per-project memory spaces	Yes; viewable/clearable in settings	Yes; default-on outside EU/UK	Yes; GA across M365 (early 2026)
Derived memory from chat history	Yes — references all past chats (Plus/Pro deeper; lighter for free)	Yes — Memory of past chats (18+, personal accounts)	Yes — synthesis updated ~every 24h; opt-in origins, broad rollout by 2026	Yes — auto-synthesized preferences, brands, dietary needs; ~95% recall claimed after Feb 2026 upgrade	Yes — persistent by default; "Skills" for reusable tasks (2026)	Yes — preferences and working style
Connected ecosystem	Connectors + in-chat apps (Expedia, Booking.com, Tripadvisor, Accor, Uber, etc.)	Deepest : Gmail, Calendar, Drive, Photos, Search/Maps/Shopping, YouTube via Personal Intelligence	Integrations/MCP tools, files, browsing; no consumer data ecosystem	Model-agnostic memory + Comet browser/assistant; enterprise connectors	X platform (social/real-time)	Deepest for work : Outlook, Teams, OneDrive, M365 graph
Where memory lives	OpenAI cloud, separate from chats	Google account activity systems	Anthropic; import/export supported (experimental)	Private per-user store; portable across models	xAI	User's Exchange mailbox , inheriting enterprise controls
Privacy/controls	View/edit/delete; Temporary Chat; memory off; enterprise exclusions from training	Per-chat toggle; activity deletion; not in EEA/UK/CH	Transparent memory use (cites when memory used); Incognito; pause/reset; enterprise admin controls	Off switch; incognito disables memory + history; encrypted	Default-on (opt-out); EU/UK excluded	Admin policies, Lockbox, encryption



Capability	ChatGPT (OpenAI)	Gemini (Google)	Claude (Anthropic)	Perplexity	Grok (xAI)	Copilot (Microsoft)
Travel-relevant strengths	Largest user base (~800M+ weekly); live OTA inventory in chat; strongest app ecosystem	Personal travel history from email/photos; Maps/Flights/Hotels data; announced agentic booking with major hotel brands	Long-context reasoning over complex itineraries/documents; transparency; strong for structured planning	Cited, real-time research; strong comparison shopping; memory works across models	Real-time events/sentiment from X; speed	Business travel: policy + calendar + email context
Key limitations	Instant Checkout for travel walked back (Mar 2026); personalization opaque to user	Ecosystem depth = privacy sensitivity; EU gap; feature fragmentation	Smaller consumer ecosystem; fewer travel integrations	Smaller scale; memory younger	Thin travel/commerce ecosystem; EU gap	Consumer travel weaker than work context

4.2 Narrative assessment

ChatGPT is the volume leader and the most important single surface for travel discovery, with independent analyses attributing the overwhelming majority of AI referral traffic to it (Conductor pegs it near 87% of AI referrals on average). Its personalization stack (saved memories + full chat-history reference + connectors/apps) is mature, and its in-chat travel apps put live Expedia/Booking inventory inside the conversation. Its March 2026 retreat from native Instant Checkout in travel — refocusing on discovery while apps handle transactions — is a telling signal about where the hard edges of agentic commerce lie.

Gemini has the most structurally powerful personalization position for leisure travel: Personal Intelligence connects the assistant to the user's actual travel record (confirmations in Gmail, trips in Photos, intent in Search and YouTube). Combined with Google's travel vertical (Flights, Hotels, Maps) and announced agentic booking with Booking.com, Expedia, Marriott, IHG, Choice, and Wyndham — where partners remain merchant of record — Google is assembling an end-to-end personalized travel pipeline. Its constraint is regulatory and reputational: the same depth that powers personalization keeps the feature out of the EEA/UK and demands careful trust management.

Claude differentiates on transparency and control: it tells users when memory informs a response, keeps project contexts separate, offers incognito chats, and supports experimental memory import/export. Its consumer travel-commerce ecosystem is thinner, but its long-context reasoning makes it strong for complex itinerary construction, document-heavy planning (group contracts, corporate policies), and analysis. (Disclosure: this report was prepared with Claude; readers should weigh platform descriptions against the cited primary sources.)

Perplexity pioneered "context portability" — one memory layer that persists while the user switches underlying models — and pairs it with citation-first answers, which matters for high-stakes travel decisions. Its memory explicitly captures brands and dietary needs, and its Comet browser/assistant pushes toward agentic task execution. Scale is the constraint.



Grok offers default-on memory and unmatched real-time social signal from X (useful for events, conditions, buzz) but has the least developed travel-commerce and connected-data ecosystem of the six.

Copilot is the sleeping giant of *business* travel personalization: memory now lives in the user's Exchange mailbox with enterprise controls, and the assistant natively sees calendars, email, and policy documents. For corporate travel programs and hotels courting business transient and group demand, Copilot-mediated discovery deserves specific attention that most hospitality GEO programs currently ignore.

5. Impact on the Customer Journey

5.1 Stage-by-stage analysis

Inspiration. This is where AI has penetrated deepest and where personalization matters most.

Surveys consistently rank discovery and recommendations as the top AI travel use case

(75% of AI-inclined travelers in IMG's 2026 survey; 65% "discovering activities and restaurants" in IPX1031's US survey). A memory-enabled assistant does not brainstorm from a blank slate: it proposes destinations that extend the user's revealed taste ("you loved Kyoto's quiet temples — consider Luang Prabang in dry season"). For destinations, this replaces the serendipity of glossy inspiration media with taste-graph adjacency — a fundamentally different marketing problem.

Destination selection. Personalization compresses comparison. Instead of the traveler weighing eight destinations, the assistant pre-weighs them against climate preference, budget, flight access from the preferred airport, school calendars, and risk tolerance, and presents two or three with reasoning. Industry observers already worry about concentration effects: well-covered destinations get recommended more, while under-documented (often independent, rural, or developing-market) options struggle to surface — a dynamic EHL and Horwath HTL analysts flagged publicly in 2026.

Hotel discovery and comparison. The AI assembles a consideration set of typically three to five properties, matched to the persistent profile, often with pros/cons framed to the user's known priorities. Cendyn's analysis puts it starkly: in a traditional results page many hotels compete for attention; in an AI answer only a few are mentioned, and often one stands out. Comparison then happens *inside the conversation* — "which has the better kids' club?" "which is quieter?" — drawing on reviews, editorial coverage, and the property's own content. A property absent from the AI's retrievable knowledge is not losing a click; it is absent from the entire deliberation.

Shortlisting and pricing research. With OTA apps inside ChatGPT and Google's AI Mode surfacing live rates, users increasingly see real pricing and availability in-conversation. Personalization shows up as default filters (the assistant checks dates from your calendar, applies your loyalty rate, knows you always want breakfast included). Price-comparison behavior persists — around half of US AI users compare flight and hotel prices with AI — but it is mediated by the assistant's framing of "good value *for you*."

Booking. Here the evidence counsels humility. Only ~13% of AI-inclined travelers in IMG's survey expect to book via AI; Skift's 2025 data found just 2% willing to give AI full booking autonomy, and McKinsey puts comfort with hands-off booking below 5%. OpenAI's walk-back of



Instant Checkout for travel (March 2026) — after which Expedia and Booking shares jumped — confirmed that travelers research in chat but transact where trust, guarantees, and servicing live. Current architecture: AI is the front door; OTAs, brand.com, and hotel booking engines are the cash register. Google's announced agentic booking keeps partners as merchant of record for the same reason. Expect assisted booking (AI fills the cart, human confirms) to normalize well before full autonomy.

Pre-arrival planning. Memory-enabled assistants become trip binders: they recall the booking, build day-by-day plans around the traveler's pace and diet, monitor for changes, and (with calendar access) slot activities around fixed commitments. Hotels that expose pre-arrival options (spa, dining, kids' programs) in machine-readable form get pulled into these plans; those that require phone calls do not.

On-property. In-destination, the assistant is a concierge that already knows the guest: "quiet breakfast place within a 10-minute walk that has gluten-free options" is answerable in one shot. This puts AI in direct competition with the hotel's own concierge and ancillary-revenue channels — unless the hotel's outlets, hours, menus, and booking links are the most retrievable answer.

Post-trip. The trip itself becomes memory. Feedback the traveler shares with the assistant ("the resort was great but the beach was crowded") permanently reweights future recommendations — an invisible, unappealable review with an audience of one very influential reader. Post-trip, assistants can proactively suggest the next trip ("you loved the Dolomites in summer; here's a winter version"), collapsing the re-inspiration cycle.

5.2 AI-personalized planning versus traditional Google Search

Dimension	Traditional Google Search	Personalized AI Assistant
User knowledge	Query + coarse signals (location, history) feeding ad targeting; results largely uniform	Deep persistent profile (preferences, history, constraints) shaping the substance of the answer
Output	Ranked links + ads; user filters	Synthesized answer naming 3–5 options; AI filters
Consideration set size	10+ visible options, pages more	Typically 3–5; often 1 highlighted
Paid influence	Direct (search ads, metasearch)	None currently on the answer itself; influence via content, data, citations
Where evaluation happens	On brand/OTA/review sites after the click	Largely inside the conversation, before any click
Traffic character	High volume, mixed intent	Low volume (≈1% of hotel site visits, mid-2026), high intent; multiple studies show substantially higher conversion than non-branded organic (hotels/resorts among the biggest gainers), though Adobe finds AI-referred travel visitors still convert 28% less than non-AI



Dimension	Traditional Google Search	Personalized AI Assistant
		sources overall while engaging far more — a gap that narrowed ~70% since late 2024
Measurability	Mature (Search Console, ads platforms)	Immature; heavy "dark" influence (research in chat, booking attributed elsewhere)
Personalization visibility	Marketer can observe/targets segments	Profile invisible to marketer; segmentation happens privately

The composite picture: search is being unbundled into *personalized deliberation* (in the AI) and *transaction* (on trusted platforms), with the deliberation phase — where brand preference is actually formed — moving somewhere marketers cannot see and cannot buy.

6. Implications for Hotels and Resorts

6.1 Marketing in a profile-mediated world

Hotel marketing has spent two decades optimizing the observable funnel: impressions, clicks, sessions, conversions. Personalized AI collapses much of that funnel into a private conversation. The strategic response is not one tactic but a reorientation around a single question: **when an AI that knows a specific traveler deliberates on their behalf, is your property the confident match?** That decomposes into visibility (does the AI know you exist and retrieve you), legibility (does it understand precisely who you're for), credibility (do independent sources corroborate you), and transactability (can it act on you — check rates, link to booking, eventually book).

6.2 SEO, AEO, and GEO

Traditional SEO remains the foundation — AI crawlers depend on indexable, well-structured sites, and a technically weak site is invisible to AI as well. On top of it:

AEO (Answer Engine Optimization) targets direct-answer surfaces (AI Overviews, snippets, voice): concise, factual answers to the questions travelers actually ask (check-in times, parking, pet policy, accessibility specifics), marked up with schema (Hotel, FAQ, Review, Offer structured data).

GEO (Generative Engine Optimization) targets the narrative layer: how AI models describe, summarize, and *recommend* the property. Practices with evidence behind them: detailed pages for each facility and audience (spa, family rooms, pet-friendly, accessible rooms) rather than generic copy; original, specific content AI can quote; consistency of facts (hours, address, amenities) across the website, Google Business Profile, and directories — inconsistency erodes machine confidence; and presence in the third-party sources AI cites. Citation-pattern research is directly actionable: analyses find Gemini draws over half its citations from brand-owned sites while ChatGPT draws nearly half from third-party directories and editorial sources, and Perplexity cites many sources per answer. Translation: your own site quality matters most for Gemini; OTA listings, Tripadvisor, and earned media matter most for ChatGPT; being cited *somewhere reputable* matters everywhere. One cautionary research finding (Wellows): a major chain earned hundreds of *implicit* AI mentions but few explicit ones — AI used its strengths to



describe categories while recommending competitors with cleaner data. Legibility is not optional.

Personalization-specific GEO is the frontier this report highlights: because AI matches user attributes to property attributes, content should be written as *machine-matchable claims about who you serve*. "Family-friendly" is weak; "kids' club for ages 4–12, open 9–5 daily, included in rate; cribs and bed rails free; two connecting-room categories; children's menu with allergy protocols" is a set of facts an assistant can confidently match to "user has children aged 6 and 9, one with a nut allergy." The same logic applies to accessibility (exact door widths, roll-in shower inventory, pool lifts), dietary programs, pet policies (fees, weight limits, relief areas), and business facilities. Specificity is the currency of personalized recommendation.

6.3 Direct booking, loyalty, CRM, and first-party data

Direct booking. AI referral traffic currently arrives largely commission-free and pre-qualified; Lighthouse reports independent hotels capturing a disproportionate share because their content is more distinctive than homogenized chain copy. The playbook: ensure the booking engine is fast, transparent, and rate-competitive on arrival (AI-referred visitors bounce if direct value isn't obvious); track AI referrers in analytics now; and treat parity seriously, because assistants increasingly *show* the price comparison.

Loyalty programs become more valuable, for a non-obvious reason: loyalty membership is one of the few durable, user-volunteered facts that travels *into* the AI's profile ("I'm Hilton Diamond") and, in agentic flows, authenticates *out* of it (the agent books the member rate). Brands should explicitly encourage members to tell their assistants about their status and preferences, publish loyalty benefits in machine-readable form, and prepare APIs for agent-mediated authentication (Google's agentic partners are already positioning here).

CRM and first-party data. The hotel's guest profile and the AI's user profile are converging on the same person from opposite sides. Hotels cannot access the AI's profile — but they can make their own so good that agents and travelers prefer to route through it: preference capture at every touchpoint, stay-history-driven offers, and (increasingly) exposing guest-authorized preference and booking endpoints so a traveler's agent can act on the relationship. McKinsey's framing is apt: the strategic goal is to be "the authoritative source of data and execution" for external AI agents.

Content strategy shifts from campaign-led to reference-led: evergreen, specific, structured, honest. Editorial coverage and reviews are content strategy too — they are what the AI reads when it doesn't trust your brochure.

Reputation management gains a second audience. Reviews are no longer only social proof for humans; they are training signal and retrieval corpus for machines, read in aggregate and thematically. Themes in reviews ("consistently praised for accessible design," "recurring complaints about noise") surface verbatim in AI comparisons. Volume, recency, response quality, and thematic consistency across Google, Tripadvisor, and OTAs all feed machine confidence.

Brand positioning must survive compression to one sentence, because that is often all the AI will say about you. A property that cannot be summarized distinctively will be summarized generically — and generic loses to specific in a personalized match. "A 41-room adults-only



cliffside inn known for its wine program and silence" is an AI-native positioning; "luxury and comfort in the heart of paradise" is noise.

6.4 Risks to manage

Concentration risk (a handful of assistants gatekeeping demand, with fee structures likely to follow the currently free traffic — several industry voices warn of an "OTA 2.0" dynamic); hallucination and misattribution risk (AI describing your property wrongly, or crediting your strengths to competitors — monitor AI answers about your property the way you monitor reviews); data-freshness risk (stale rates, closed outlets, outdated policies quoted confidently); equity risk for small and developing-market operators with thin digital footprints; and privacy/brand-safety risk in how hotels themselves use guest data to meet the personalization expectations AI is setting.

7. Future Outlook: 2026–2030

Forecasting in this domain requires discipline about confidence levels. We divide the outlook into **likely developments** (strong current momentum, already shipping in early form), **emerging trends** (real but unproven at scale), and **speculative possibilities** (plausible, contested, dependent on trust or regulation breakthroughs).

7.1 Likely developments (high confidence, 2026–2028)

Persistent AI companions become the default travel interface for research. Memory shifted from differentiator to baseline across all major platforms during 2025–26. By 2028, expect most trip planning by AI-using travelers to begin with an assistant that already knows their travel identity, with the "cold start" query becoming rare.

Assisted (human-confirmed) agentic booking normalizes. The pattern established in 2025–26 — AI assembles, human confirms, supplier transacts as merchant of record — scales. Google's AI Mode agentic booking (announced with Booking.com, Expedia, Marriott, IHG, Choice, Wyndham) and ChatGPT's app-mediated commerce are the templates. Fully autonomous booking remains niche through at least 2028 given 2–5% current comfort levels, but comfort will climb as agents prove reliability on low-stakes tasks (restaurant reservations, tickets — already live in Google's US rollout).

Calendar- and email-aware recommendations spread beyond Google/Microsoft.

Connected-context personalization is too valuable to remain proprietary; expect standardized, user-authorized context sharing (the Model Context Protocol pattern, already adopted in hospitality tech by Apaleo and Sabre) to let travelers grant any assistant scoped access to bookings and calendars.

Loyalty integration. Agent flows that authenticate loyalty membership and book member rates move from pilots to standard, making programs a primary AI-era moat for brands that execute.

Real-time itinerary optimization. Agents that monitor flights, weather, and closures and proactively rebook/resequence — early versions live in airline concierges (e.g., Delta's) and OTA assistants (Expedia's Romie) — become table stakes for premium travel products.

Voice and multimodal planning grow steadily. Voice-first planning (in cars, via earbuds and smart glasses) and photo-based inspiration ("find me somewhere that looks like this") expand,



favoring properties with strong visual assets and clean structured data that survive audio compression to a one-sentence recommendation.

7.2 Emerging trends (medium confidence, 2027–2029)

The personal travel agent model. Assistants graduate from answering to *managing*: standing preferences ("keep me under \$400/night, aisle seats, no red-eyes"), budget envelopes, and proactive proposals timed to school holidays and fare dips. McKinsey/Skift's "direct-booking renaissance" scenario — supplier AI negotiating tailored offers directly with traveler agents — becomes visible in luxury and corporate segments first.

Wearables and ambient context. Health data (sleep, activity, HR) informing wellness-travel matching and in-trip pacing; location-aware on-property suggestions. Adoption depends on comfort with intimate data sharing and will vary sharply by market.

Memory portability and user-owned profiles. Perplexity's cross-model memory and Claude's memory import/export point toward traveler-owned preference profiles that move between assistants — which would commoditize the assistants' data moats and raise the value of supplier-side data quality.

AI-to-AI commerce plumbing. Standardized machine-readable offers (dynamic pricing, cancellation rules, ancillaries) and robust booking/servicing APIs — which McKinsey identifies as preconditions for agentic scale — become a competitive requirement for hotel technology stacks, alongside new commercial agreements on attribution, liability, and revenue share.

Monetization of AI travel surfaces. Currently commission-free AI referrals are unlikely to stay that way; expect sponsored placement debates, agent-facing rate deals, and possibly per-transaction fees. Hotels should bank the free-traffic era while it lasts and scenario-plan for its end.

7.3 Speculative possibilities (low confidence, 2028–2030)

"Agent takeover" end-state. McKinsey/Skift's scenario in which customer agents and supplier agents negotiate and transact autonomously, compressing OTAs and metasearch. The 2026 evidence (Instant Checkout's travel walk-back; entrenched trust in booking platforms) suggests this is possible but far from inevitable; travel's complexity — volatile pricing, servicing, liability — is a genuine moat for incumbents.

Fully autonomous booking as mainstream behavior. Requires the trust curve to move from ~5% to ~50%; contingent on flawless agent performance, clear liability frameworks, and probably regulation (EU AI Act transparency obligations, payment-authentication rules) maturing around agent transactions.

Hyper-personalized dynamic packaging and pricing at the individual level, negotiated agent-to-agent — raising fairness and regulatory questions (personalized pricing is already a sensitive area under EU consumer law) that could cap how far this goes.

Privacy regime shocks. GDPR-style constraints on derived profiles could expand (the EEA/UK exclusions of 2025–26 are the leading indicator), potentially bifurcating the world into deep-personalization and shallow-personalization markets — with direct consequences for how hotels prioritize GEO investment by source market.



8. Strategic Recommendations by Stakeholder

Independent hotels. You are, counterintuitively, advantaged: Lighthouse data shows independents capturing a disproportionate share of AI referrals because distinctive content beats homogenized chain copy, and AI recommendations can route commission-free demand straight to your booking engine. Priorities: (1) fix the data layer — consistent NAP/amenity facts everywhere, full schema markup, fast indexable site; (2) publish audience-specific specificity (families, accessibility, pets, dietary) as verifiable facts; (3) make direct booking obviously easy and price-transparent for AI-referred visitors; (4) monitor AI answers about your property monthly across ChatGPT, Gemini, Perplexity, and Copilot and correct source data driving errors; (5) invest in reviews and a handful of earned-media citations per year.

Luxury hotel brands. Personalized AI rewards sharply articulated identity and punishes generic "5-star" language. Codify what the brand is *for whom* in one-sentence, machine-quotable form per property. Prepare for agent-mediated clienteling: guest preference APIs, recognition across stays, and pre-arrival personalization that meets the standard AI is setting. Protect the human layer — as AI absorbs logistics, human connection becomes the scarce luxury good; market it explicitly. Watch wearable/wellness data integration as an early differentiation arena.

Destination marketing organizations. Your battleground is the inspiration stage, where AI influence is strongest. Fund the destination knowledge layer: comprehensive, factual, structured content on seasons, regions, accessibility, and experiences that AI can retrieve and trust — including for lesser-known areas, to counter AI's concentration bias toward well-documented places (and to spread demand away from over-touristed icons). Cultivate the third-party citation ecosystem (guidebooks, editorial, credible blogs). Measure AI share of voice for the destination against competitors as a core KPI.

Travel marketers and hospitality agencies. Build GEO/AEO as a service line with measurement at its core: AI referral tracking (regex-based channel grouping today), answer monitoring, citation analysis by platform (own-site emphasis for Gemini; directories/editorial for ChatGPT). Reweight budgets gradually from pure paid-search toward content, data, PR, and review infrastructure — the inputs AI actually consumes. Educate clients that AI influence is mostly "dark" (research in chat, conversion attributed elsewhere) and set expectations accordingly.

Revenue managers. Prepare for demand that arrives pre-filtered and better informed: AI-referred guests have often already compared your rates. Rate parity and value framing (what's included) become machine-read, so errors are amplified. Begin scenario work on agent-era distribution: machine-readable offers, personalized/member rates surfaced via API, and the possibility of per-referral or per-transaction AI channel costs. Track AI-referred booking cohorts separately — early data suggests different lead times, lengths of stay, and cancellation behavior are plausible and worth quantifying on your own book.

Digital marketers (in-house). Near-term checklist: implement AI-source analytics segmentation; audit the property in each major assistant quarterly (generic prompts *and* persona-style prompts — "family with young kids," "wheelchair user," "traveling with a dog"); close every factual inconsistency the audits reveal; add FAQ and amenity schema; create or upgrade audience-specific pages; brief the PR function that citations are now a performance channel; and claim/optimize presence on the OTA and review platforms that dominate ChatGPT citations.



Hotel technology providers. The build list: MCP-style agent interfaces to PMS/CRS/booking engines; AI-visibility monitoring and GEO tooling (an emerging category — Operto's free GEO Consultant signals commoditization pressure, so differentiate on action, not just reporting); guest-profile APIs with consent management for agent-mediated personalization; and content/schema automation for multi-property groups. Standards participation (MCP, agentic commerce protocols) is strategic, not optional.

9. Evidence Base, Limitations, and Open Questions

This report distinguishes documented fact from analysis. **Well documented:** platform memory mechanics (vendor documentation from OpenAI, Google, Anthropic, Microsoft, Perplexity, xAI); AI adoption ranges in travel planning (Kantar/Statista ~40% globally; IMG 33% likely in 2026; IPX1031 ~30% of Americans; Klook 91% among its global user base; McKinsey/Skift >50% of US travelers — definitional differences explain the spread); the discovery-heavy, booking-light usage pattern; AI referral traffic scale (~1% of hotel site visits) and rapid growth (Lighthouse); higher AI-referral conversion versus non-branded organic in multiple independent analyses, alongside Adobe's finding of a persisting (but fast-narrowing) overall conversion gap; platform citation patterns; and the 2025–26 agentic-commerce events (ChatGPT travel apps, Instant Checkout retreat, Google agentic booking partners).

Informed analysis, clearly flagged as such: the signal-impact ranking in Section 3 (no platform publishes weights); the persona illustration in Section 2; the degree to which persistent profiles currently alter recommendation sets versus merely framing (systematic A/B evidence across real user profiles is thin — platforms don't expose it, and academic study is nascent); and all 2027–2030 projections.

Evolving/limited evidence: attribution of AI-influenced bookings (most AI influence is invisible to analytics); how memory-driven personalization behaves for European users under divergent regulation; whether AI concentrates or diversifies destination demand at scale; the durability of commission-free AI referrals; and the true trajectory of booking autonomy, where surveys, platform behavior, and vendor marketing tell conflicting stories. Hoteliers should treat vendor-published statistics (including several cited here from companies selling GEO/AI tools) with appropriate caution and prioritize measurement on their own traffic.

10. Key Takeaways for Hotel Marketers

- **The consideration set is now assembled by a machine that knows the traveler.** Winning means being the confident match for specific, known people — not the loudest generic message.
- **You can't see or buy the profile; you can only be matchable against it.** Specific, structured, verifiable content about *who you serve* (families by age, accessibility in measurements, pets in policies, diets in protocols) is the highest-ROI personalization response available today.
- **Consistency is credibility.** AI cross-checks your website, Google profile, OTAs, and reviews; contradictions cost recommendations.



- **Different assistants read different sources.** Own-site depth for Gemini; directories, OTAs, and editorial for ChatGPT; citations everywhere for Perplexity; don't ignore Copilot for business travel.
- **AI traffic is small, fast-growing, high-intent, and (for now) commission-free.** Instrument it, convert it, and bank the free era.
- **Loyalty and first-party data are your seat at the agentic table.** Encourage guests to teach their assistants who they are to you, and build the APIs agents will call.
- **Discovery is being won and lost in chat; booking still happens on trusted rails.** Optimize for the conversation without abandoning the booking engine, parity, and post-booking excellence that still close the deal.
- **Audit yourself in the machines quarterly — including through personas.** What the AI says about you to a family, a wheelchair user, and a honeymooner is your new brand tracker.

Selected Sources

Platform documentation: OpenAI Memory FAQ & memory announcements (help.openai.com; openai.com); Google Gemini Personal Intelligence & memory help pages (gemini.google; support.google.com); Anthropic Claude memory documentation (support.claude.com); Perplexity, "Introducing AI assistants with memory" (perplexity.ai); Microsoft 365 Copilot Memory documentation; xAI Grok memory announcements.

Industry research and press: McKinsey & Skift, "Remapping Travel with Agentic AI" (2025) and McKinsey Live follow-ups (2025–26); Skift coverage of OpenAI ChatGPT apps and Instant Checkout walk-back (2025–26); PhocusWire coverage of ChatGPT travel apps and agentic commerce; Lighthouse AI-referral data via Hospitality Net/eHotelier (2026); Adobe Digital Insights travel AI-readability and conversion data (2026); Seer Interactive, First Page Sage, Visibility Labs, Conductor, and BrightEdge AI-traffic conversion/share analyses (2025–26); Kantar via Statista, AI use in travel planning (2024–25); IMG 6th Annual Travel Outlook Survey (2026); IPX1031 Americans' Travel Report (2026); TakeUp "Rise of AI-Planned Travel" (2026); Klook global traveler survey via CNBC (2026); Accenture consumer AI-travel research via PhocusWire; Hotel Dive on Operto GEO Consultant (2026); Cendyn, UP Hotel Agency, HiJiffy, GuestRevu, Hotelrank on GEO/AEO for hotels (2025–26); Google AI Mode agentic booking coverage (2025–26); Expedia "Expedia in ChatGPT" (2025–26).

Figures such as "91% of travelers use AI planners" (Klook) reflect specific surveyed populations and should not be generalized; ranges across surveys are presented deliberately.

Prepared July 2026. This report synthesizes publicly available sources current to early July 2026; the platform landscape described here changes monthly and material facts should be re-verified before major investment decisions. Nothing herein is legal advice regarding privacy regulation.

