

Beyond the Booking Funnel

AI in the In-Destination Layer: Museums, Heritage, Tours, Restaurants, Attractions and Car Hire (2022–2026)

Version 1.0 — September 2026 (first release)

Note on this version. This is the first release, preceded by two internal drafts (June 2026 and a September 2026 update). The September draft was submitted to an independent evidence audit before release; six statements found to be misstated or overstated are corrected here and logged in Section 10. Every statistic, date, benchmark result and legal proposition carries a numbered citation to Section 11, where sources are labelled by evidence type. Qualitative judgements are not shown as measurements, unlike denominators are not plotted on a shared axis, vendor-reported results are described as such, and interpretation is distinguished from disclosed fact.

Scope and argument. The June 2026 draft argued that the in-destination layer runs on a different logic from the booking layer: in flights and hotels the central AI question is who captures the transaction, whereas in museums, tours, restaurants and attractions it is who shapes *discovery and the on-site experience*. Three months on, that argument is consistent with the new evidence, which is now more European. What has changed is who is competing to shape discovery: Google, OpenAI, Anthropic and the experience marketplaces have all moved into in-destination discovery since June, and the destination marketing organisation (DMO) has become one of the most visibly contested segments in the evidence reviewed. This version adds a section on the destination layer and a focus section on the agentic harnesses now in production.

What has changed since June 2026 — seven developments

1. Planning use is high; observed AI-referred bookings are small — on non-comparable indicators.

Vendor research previewed by GetYourGuide reports 69% of surveyed travellers have used AI to plan or book a trip and 17% to book an activity [13]; Booking Holdings disclosed that LLM-referred traffic is well under 1% of room nights [19].

2. Europe has its own data. Phocuswright's Europe Consumer Travel Report 2026 (UK, France, Germany) finds in-destination AI use is led by real-time recommendations, and that European travellers use AI more while trusting it less than human sources [1,2,3].

3. Big platforms entered in-destination discovery. Google's Ask Maps gained event discovery (August) [6]; Viator became Google's first connected app for experiences (July) [15]; Tripadvisor and Viator launched connectors in Claude (April) [16]; OpenAI moved ChatGPT checkout to partner apps (March) [17,18].

4. DMOs became a battleground. Mindtrip entered Europe in June and has announced destination partners at intervals of roughly three weeks since, including a WhatsApp channel [4]; its commissioned research reports 22% of surveyed DMOs offer AI trip planning [5].

5. Heritage deployments got cheaper and more contested. London Transport Museum reports a conversational exhibit delivered for under £10,000 [8]; the Canadian Museum of History's chatbot drew public objections about replacing guides [10].

6. Regulation arrived. EU AI Act Article 50 transparency duties applied from 2 August 2026; high-risk obligations were deferred by the Digital Omnibus [12].

7. Chatbots became harnesses. Priceline relaunched Penny as a system of more than ten agents [20]; Sabre reports a production MCP deployment [21]; Amadeus announced Claude-oriented tooling [22]; the payment networks announced agent mandates [23,24,25]; benchmarks moved to groups and verifiability [27,28].

1. A Different Logic: Discovery and Experience, Not Bypass

The disintermediation anxiety that drives the booking-layer debate largely does not transfer to the in-destination layer, and the reason is structural. A flight or a hotel room is a commodity booking whose intermediary value — option-finding — is what a conversational agent reproduces. A museum visit, a guided walk or a meal is an experience consumed on site, where AI's most active role is enriching the encounter rather than removing a middleman from the purchase.

A definition is needed, because the term does work throughout. *Disintermediation* here means the loss by an incumbent intermediary of one or more of five things: control of discovery, ownership of the transaction, ownership of the customer relationship and data, attribution, and margin. A supplier can keep fulfilling — running the tour, serving the meal, opening the gallery — while an AI interface takes over some or all of those five. Continued physical fulfilment is therefore not, by itself, evidence that a venue is unaffected; the question in each segment is which of the five are moving, and toward whom.

The summer's evidence is consistent with this. Phocuswright's European survey reports that real-time recommendations are the leading in-destination use of AI in each market studied, reaching 53% of AI-using UK travellers, with operational tasks such as disruption handling and reservation management trailing by twenty points or more [2]. This is survey evidence of reported use, not a causal test; but it is independent of vendors and it points the same way as the supply-side indicators in Section 4.

In the in-destination layer, AI remains a tool that institutions and operators adopt to deepen engagement rather than a force that bypasses them. What is new is that the

largest technology platforms have decided the discovery layer is worth contesting directly.

Figure 1 restates the June draft's map of where AI activity falls in each segment as a qualitative matrix. The ratings are the author's judgement from the deployments, product launches and surveys cited in Sections 2 to 7; they are not measurements, and the earlier stacked-bar presentation, which implied shares, has been withdrawn.

Segment	Discovery & engagement	Booking & transaction	Back-office & operations	Basis (see sections)
Museums & heritage	High	Low	Low–medium	Guides, visitor Q&A, translation; ticketing largely untouched (§3)
Destination (DMO) layer	High	Low	Low	Trip planners, WhatsApp channels, Ask Maps; DMOs rarely transact (§2)
Guided tours & activities	Medium–high	Medium	Low–medium	Marketplace assistant integrations; booking completes on marketplace (§4)
Restaurants	Medium	Medium–high	Low–medium	Network-wide Concierge; reservations remain platform-mediated (§5)
Attractions & theme parks	Medium	Medium–high	Low	Bundled into OTA Connected Trip; Ask Maps event links (§6)
Rental cars	Low	Medium	High	No new evidence found in the search scope (§7)

Figure 1. Where AI acts within each in-destination segment: qualitative matrix (author's judgement, low / medium / high). Criteria: 'High' = the role is the primary stated purpose of most cited deployments; 'Medium' = present in several; 'Low' = rare or absent in the evidence reviewed.

2. The Destination Layer

The June draft treated destinations only implicitly. Since June the destination marketing organisation has become the segment where AI vendors, platforms and consumers are moving fastest, and where the competitive dynamics look least like the museum case and most like the booking-layer land-grab.

2.1 Vendor entry: the Mindtrip cadence

Mindtrip, a San Francisco conversational-planning platform founded in 2023, announced its European launch on 3 June 2026 with first contracts at Madeira Promotion Bureau and Norwegian Travel Cluster, framed as a response to declining organic search [4]. Company releases since then name Hilton Head (10 June), Savannah and Tybee Island (16 June), Meet Boston (30 June), San Diego Tourism Authority (28 August) and Meet Minneapolis (2 September), on top of an existing roster including Visit California, Discover Puerto Rico and San Francisco Travel [4] — seven announcements between 29 April and 2 September, roughly one every three weeks, with none in July. Announcement cadence is not contract cadence and none discloses contract value. Boston is described as the first destination to launch Mindtrip's WhatsApp integration, timed to the FIFA World Cup [4]. These are vendor announcements: none discloses usage, accuracy or conversion. The cadence nonetheless matters: messaging-based destination assistance, which some smaller products have treated as a differentiator, is now offered by a well-funded vendor in the segment.


2.2 Platform entry: Google's Ask Maps

Ask Maps, the Gemini-powered conversational layer in Google Maps unveiled in March, was updated on 6 August 2026 with hotel and local-event discovery and extended to more than 150 countries and territories; the hotel, event and dining features are US-only at launch [6]. Google describes the update as handling multi-step tasks with agentic capabilities [6]. Separately, Google's AI Mode gained hotel booking on 27 August, in US English, with Booking.com, Expedia, Marriott and Hilton among launch partners; the hotel or travel partner remains merchant of record [7]. That is a booking-layer move, but it shows the direction for attractions and events. Together these place the dominant map inside the in-destination discovery moment that the June draft identified as the layer's centre of gravity. Metasearch is moving the same way: Skyscanner's 'Explore with AI' beta (summer 2026) delivers paid destination content from DMOs inside AI-generated results, another channel in which a destination's message reaches the traveller through an intermediary's assistant [31].

2.3 Two surveys of demand

The first survey is vendor-commissioned and should be read as such. Mindtrip, with Sabre, Sojern and MMGY Travel Intelligence, surveyed 1,006 US travellers and 103 DMO leaders in April 2026 and released results on 20 August [5]. It reports that 54% of travellers have used AI to plan travel and 63% would use an AI planner on a destination website, while 22% of DMO leaders say their organisation offers one and 26% have a documented AI strategy; 25% of travellers say they are less likely to use a destination site lacking AI features, and 72% of DMO leaders say non-adopters risk irrelevance within two years [5]. The last figure is stated opinion, not an observed outcome. Question wording is not published. The two samples are separate populations answering separate items, so they should not be read as a single measured gap (Figure 3).

Two survey snapshots: traveller intent and DMO provision (separate samples, separate questions)



Source: Mindtrip with Sabre, Sojern and MMGY Travel Intelligence, AI readiness study; released 20 Aug 2026. Vendor-commissioned (Mindtrip sells DMO AI tools); exact question wording not published. The panels are not subtractable: they measure different populations and items. Panel B's 72% is executive opinion, not an observed outcome.

Figure 3. Two survey snapshots from the Mindtrip / MMGY study [5]: traveller intent (Panel A) and DMO provision (Panel B). Separate samples and items; vendor-commissioned; U.S. only.

The second survey is independent of vendors. Phocuswright's Europe Consumer Travel Report 2026 was fielded in April among leisure travellers in the UK, France and Germany and published in June 2026, with AI-specific findings released on 1 and 4 September to coincide with its London summit [1,2,3]. The full report is a paid product and question wording is not in the public releases; every figure below is taken, with

attribution, from the publisher's own public releases of 1 and 4 September, not from the report itself. Three findings matter here. First, AI is shaping decisions: close to nine in ten AI-using travellers say AI influenced an actual trip decision (88% in Germany, 92% in France), and UK travellers are the most likely to say it shaped where they went [3]. Second, AI-assisted research is fragmented across channels: standalone AI platforms are the entry point for 44–51% of AI-using travellers depending on market, AI embedded in OTAs and metasearch for 21–33%, AI answers inside search engines for 25–34%, and direct supplier sites and apps for 11–16% [2]. Third, Europeans report using AI more while trusting it less than friends, family and human advisers [3], consistent with the wider conversational-agent literature, in which consumer trust is a dominant research theme [37]. The implication the author draws is that the scarce asset is not AI capability but a trusted source for the AI to draw on — which is, in principle, what a destination or heritage institution is.

For a destination, the AI question is becoming one of visibility and trust rather than technology: will the visitor's assistant surface the destination's own knowledge, or a platform's aggregated substitute?

3. Museums and Cultural Heritage (Centrepiece)

Cultural heritage remains the clearest case of AI as augmentation rather than disintermediation. The personalisation paradox the June draft described is unchanged, and conversational AI is still positioned as the bridge. The summer added evidence on cost, on public reception, on grounding, and on regulation.

3.1 Deployments since June

The most instructive new case is small. In June 2026 the London Transport Museum presented, at the Museum AI Summit 2026, a conversational guide built around a 1927 poster in its Global Poster Gallery, conversing in some sixty languages, which the museum's team described as delivered for under £10,000 and as a replicable pattern that keeps interpretive knowledge in-house [8]. The cost figure is the museum's own, from a trade case study, not an audited comparison; but if broadly representative it materially lowers the entry cost for mid-sized institutions.

The Canadian Museum of History's MUSEO is the cautionary case, and also a well-documented example of grounding. MUSEO is a QR-code-triggered assistant in the Canadian History Hall, built with the French vendor Ask Mona, answering questions in English or French about ten objects in Gallery 3; it runs on a closed model drawing only from museum-validated content and underwent seven months of internal testing [9]. The pilot was described publicly in early 2026 and in a vendor case note on 10 July [9]. In August, a museum social-media post promoting it drew hundreds of comments, many objecting that it could replace guides, and the museum stated publicly that it was a supplement, not a substitute [10]. Whether this is the first such reaction is not something this paper can establish; it is a clear instance, and it argues for launching such tools with interpretive staff visibly in the frame rather than as a cost story.

The large-institution deployments described in the June draft — the Louvre's 'Leonardo' assistant, the British Museum's AR app and chatbot, the Metropolitan Museum's 45,000-artwork chatbot, and the Museum of Tomorrow's IRIS+ accessibility features — remain the reference points at the large end, though they rest on trade and vendor reporting rather than institutional records [32,33,34]. Vendors pitch the commercial case as engagement that lifts secondary spend (dwell time, repeat visits, shop and membership), a self-interested framing whose underlying mechanism is nonetheless plausible [38]. The pattern that has emerged in 2026 is bifurcation: a few large institutions building bespoke systems, and a

growing tail adopting subscription vendor guides whose economics shift spend from capital to operating budgets and whose pitch leads with language coverage and temporary-exhibition agility rather than engagement statistics.

3.2 Accuracy: what one controlled study does and does not show

The evidentiary caution of the June draft stands: engagement percentages in vendor material are promotional and the peer-reviewed base is thin. Two peer-reviewed studies frame it. Trichopoulos, Ordoumpozanis and Caridakis (ACM JOCCH, January 2026) ran three field tests of LLM-based guides — a digital art exhibition, a painting gallery and the Archaeological Museum of Mytilene — and found the LLM guide outscoring a static chatbot on engagement and meaningful experience, while naming lack of context-awareness and hallucination as the two unresolved problems [35]. Derda and Predescu (2025) surveyed museum practitioners and proposed a human-centric governance model built around visitor-centred design, preservation of curatorial vision and ethical AI governance [36]. One further study is worth stating precisely. Dibitonto, Ferrato, Limongelli and Patroni, in *Multimedia Systems* (2026), used GPT-4 with Falk and Dierking's visitor-type framework to generate personalised and general audio-guide texts for three artworks, evaluated them with 56 self-identified participants, and had an expert review the outputs; personalised guides outperformed general ones on most needs, and the expert review found the outputs factually sound while identifying gaps in domain terminology and tone [11]. The guides were generated from the model's general knowledge of well-known works, not from a retrieval-grounded institutional corpus; the authors propose a curator-governed production framework as the next step [11]. The study therefore supports the value of personalisation and shows that a frontier model can be factually adequate on famous works under expert review; it does not validate a grounded production guide, and it says nothing about less-documented collections, where risk may be higher because the model's own knowledge is thinner. MUSEO's closed, museum-only knowledge base [9] is the production pattern that addresses that gap.

3.3 Regulation: Article 50 is now live

Any EU-facing visitor chatbot now carries a duty that did not exist when the June draft was written. The Digital Omnibus on AI, Regulation (EU) 2026/1744, in force from 27 July 2026, deferred the AI Act's high-risk obligations to December 2027 and August 2028 but left the Article 50 transparency duties on their original date of 2 August 2026 [12]. In outline: providers of systems intended to interact with people must ensure users are informed they are dealing with an AI, unless that is obvious to a reasonably informed person; providers of generative systems must mark synthetic output in a machine-readable way, with a grace period for systems already on the market to 2 December 2026; and which duties apply depends on whether an organisation is the provider or the deployer of the system [12]. Fines for Article 50 breaches run to €15 million or 3% of worldwide turnover under Article 99 [12]. For a museum or destination chatbot the notice duty is easily met, but it must be met, and a deployer using a vendor's system should confirm contractually who carries which obligation. The Omnibus also created an EU-level regulatory sandbox with priority access for SMEs [12]. This is informational analysis, not legal advice.

4. Guided Tours and Activities

This is the segment where the summer produced the most evidence, and it points the way the June draft anticipated: AI's role in tours and activities is concentrated at the discovery step, the marketplaces are integrating into assistants rather than being displaced by them. Direct measurement of native in-assistant booking remains limited: most documented tours-and-activities integrations hand off to a marketplace,

while Google's new hotel-booking flow in AI Mode is a notable counterexample. The evidence is largely vendor-published, and its bases differ, so Figure 2 presents it as three source cards rather than one chart. They are indicators of AI involvement at different points, not a funnel.

GetYourGuide's Spring 2026 release (22 April) previewed research with Arival, published in May, reporting that one in three trip-planning sessions begins with AI, that 69% of surveyed travellers have used AI to plan or book a trip, and that 17% have used it to book an activity [13]; sample design and question wording are not publicly disclosed. The company's product response was framed around that gap: AI review summaries and personalised search for travellers, dashboards and AI-suggested replies for operators [13]. The marketplaces themselves remain robust — GetYourGuide reported a record 10 million experiences booked in Q3 2025 and reached adjusted-EBITDA profitability [39] — and the long tail of small operators remains technologically immature, which keeps aggregation valuable [39]. Its Spring 2026 Travel Experience Trend Tracker, a survey of 505 operators, reports that 56% find AI overwhelming — a measure of sentiment, not capability [14]. Tripadvisor reported in June that more than half of Viator operators use its AI listing builder and claims that AI-built listings convert better, without publishing an effect size or method [15].

Source	Figure	Exact measure	Population / base	Period	Evidence type
GetYourGuide press release previewing Arival research [13]	33%	Share of trip-planning sessions said to begin with AI	Consumer survey; sample and geography not publicly disclosed	Apr–May 2026	Vendor release
Same [13]	69% / 17%	Respondents who have used AI to plan or book a trip / to book an activity	As above	Apr–May 2026	Vendor release
Booking Holdings Q2 2026 earnings call [19]	<<1%	Room nights attributable to LLM referral, described as 'significantly less than 1%'	Booking.com room nights, group disclosure	Q2 2026	Earnings disclosure

Figure 2. Three non-comparable indicators of AI use across planning and booking. They come from different populations, products, periods and denominators and are not plotted on a common axis. They suggest different levels of AI involvement at different points in the journey; none observes a common funnel or the rate of native in-assistant transactions, so they do not establish a conversion gap.

The distribution picture changed materially. Tripadvisor and Viator launched connectors in Anthropic's Claude in April (Viator for experiences; Tripadvisor for hotels) and expanded voice search in Amazon's Alexa+ [16]; on 30 July Viator became Google's first connected app for travel experiences, exposing more than 425,000 experiences inside Gemini, with booking completed on Viator [15]. In each case the booking still completes on the marketplace's own site.

OpenAI's checkout history is worth stating accurately because it is often misreported. OpenAI launched Instant Checkout in September 2025 for retail merchants (Etsy first, Shopify merchants to follow), with the user confirming order and payment details in chat, and separately launched apps in ChatGPT in October 2025, through which Expedia, Booking.com and others surface live inventory [17]. In March 2026 an OpenAI spokesperson said Instant Checkout was moving to apps, with ChatGPT prioritising search and discovery and purchases completing within connected services; travel shares rose on the news, and Skift noted that OpenAI had not explicitly targeted travel transactions [18]. No primary record reviewed here shows a native travel checkout product having launched; the pre-release draft's statement that one was launched and withdrawn was wrong and is corrected. Booking Holdings then put a number on assistant-referred volume: LLM-referred traffic remains significantly under 1% of room nights, even as

management reports growing visibility in AI answers and cites Google's AI Overviews as a search headwind [19]. That metric is referred room nights, not all AI-influenced bookings.

The strategic question for the marketplaces has therefore shifted from whether the assistant will replace them to which assistants will surface them and on what terms — a visibility contest rather than a disintermediation contest. For small operators the practical advice in the trade is blunt: no individual operator will build a ChatGPT app; the lever is to be listed on the marketplaces that have one and to structure one's own web content so that AI tools surface it at the research stage.

5. Restaurants

The June draft described restaurant reservation as the in-destination segment where a new intermediary most clearly captures the layer. That has intensified along the discovery axis. OpenTable launched Concierge on individual restaurant profiles in 2025; on 14 July 2026 it announced that Concierge had moved to its homepage, turning a per-restaurant Q&A tool into network-wide search across more than 65,000 restaurants, built with Perplexity and OpenAI [29]. Trade coverage read this as an attempt to intercept the 'where should I eat' question before it reaches a search engine. OpenTable is simultaneously integrated with ChatGPT, Gemini, Perplexity, Copilot and Alexa+ [29].

Two details matter for adjacent builders. First, OpenTable has no open public API: developer documentation exists, but production access to reservation and guest data is application-, review- and partner-gated [29], so the platform rather than the venue decides which assistants get in; the restaurant's own economics are a per-cover fee of US\$1.00–1.50 on network bookings plus a US\$149–499 monthly plan, rather than a percentage commission [40]. Second, OpenTable's consumer research released with the 14 July announcement reports that 61% of respondents have chosen a destination because of its food or restaurant scene, and 51% plan to use AI tools to discover and book restaurants while travelling this year [30]. The methodology is published: an online survey by Ripple Research of 1,500 US residents who had dined in a hotel restaurant within the past five years, fielded 28 May to 2 June 2026. That sample is hotel-restaurant diners, not Americans at large, so OpenTable's own '61% of Americans' framing overreaches; within its population the figure underlines how far dining has moved toward being a discovery driver. Figure 1 rates restaurants 'medium' on discovery for that reason.

6. Attractions and Theme Parks

Attractions remain largely inside the bundling strategies of the large platforms. Booking Holdings' second-quarter results (4 August) disclosed attraction-ticket growth in double digits and Connected Trip transactions growing in the low double digits, more than twice the rate of overall transactions and now a low-double-digit share of Booking.com's total [19]. The deceleration from the near-80% attraction growth reported for full-year 2025 [41] is consistent with a launch-phase line maturing rather than a retreat; that the segment is being absorbed into the bundle is the author's interpretation, not a disclosed conclusion. Google's Ask Maps event discovery, with links to tickets, is a sign of a platform contesting the attraction-discovery moment outside the OTA bundle [6], and agentic booking of attractions inside AI Mode would be the natural extension of the hotel launch [7].

7. Rental Cars

The author found no new evidence of a distinct AI-driven shift in how rental cars are discovered or booked. Search scope: English-language trade press (Skift, PhocusWire, Travel Weekly), company press releases and earnings calls for Booking Holdings and Expedia, June to 7 September 2026. This is a statement about what was found within that scope, not a market-wide finding; the segment remains operations-led on the supplier side (dynamic pricing, fleet management, automated inspection) [42] and bundled on the customer side, and is retained for completeness.

8. Focus: Custom Agentic Harnesses and Multi-Agent Workflows

The June draft described AI in travel largely as chatbots. The systems announced since spring 2026 are better described as *harnesses*: an orchestrating model wrapped in tools, memory, inventory connectors, guardrails and payment mandates, often with specialised agents beneath it. This section looks at what has been announced or deployed, what protocols it runs on, what the benchmarks say, and what follows for an in-destination harness. Most of the production evidence is company-reported.

8.1 Production multi-agent systems: Priceline's Penny as the reference case

The clearest public account of a production travel harness is Priceline's relaunch of Penny on 3 June 2026. Priceline describes Penny as a coordinated system of more than ten specialised agents covering hotel search and recommendation, flights, rental cars and customer service, on a proprietary stack combining preference learning, memory and deals technology with third-party models: Anthropic's Claude for core conversational reasoning and planning, and OpenAI and Google Cloud components for search and voice [20]. A request comparing flights to three cities for a given week is decomposed, dispatched and returned with maps and live inventory, and the traveller can book in the conversation [20]. Priceline reports that engaged users show higher conversion and that Penny usage has reduced support contacts, without effect sizes or independent validation [20]. Its chief technology officer's framing is the one worth retaining: the advantage lies not in the model but in connecting it to context, inventory and deals [20]. Booking.com's AI trip planner, Agoda's visual search and the group-wide voice rollout described on the Q2 call are variations on the same architecture [19].

8.2 The plumbing: MCP as the industry's tool layer

The less visible development is the convergence of travel-technology incumbents on the Model Context Protocol as the way agents reach inventory. Sabre launched agentic APIs behind a proprietary MCP server in September 2025, describing it as a universal translator for any AI agent, with candidate use cases including a call-centre proxy agent, a hotel-operations agent, a visa and compliance agent and an expense-filing agent [21]. In June 2026 it reported a production deployment: Linex Travel (the former Ultra Group) runs agentic workflows on Sabre's MCP server for ticket reissues, exchanges and servicing updates across a loyalty-redemption marketplace in twelve countries, with further customers said to be piloting [21]. These are Sabre's statements; no independent performance data has been published. Amadeus, having closed its self-service developer portal in July, announced on 3 September a collaboration with Anthropic: a developer toolkit that would expose flight and destination content, connectivity, documentation and authentication inside Claude Code, and a Claude Cowork plugin for travel professionals [22]. At the time of writing both are announced rather than available — the toolkit is wait-listed and the plugin under development, and access terms and pricing have not been published [22]. OpenAI's apps in ChatGPT, through which Expedia, Booking.com and Tripadvisor surface inventory, are built on MCP [17]. The practical consequence, if these initiatives mature, is that a custom harness would need an MCP client and a commercial relationship rather than bespoke integrations. The commercial relationship is

the hard part: Amadeus terms are unpublished, Sabre's server sits inside enterprise contracts, and OpenTable has no open public API [21,22,29]. The tool layer is standardising technically while remaining closed commercially.

8.3 The last mile: payment mandates and checkout

Two things happened at the transaction end. OpenAI moved ChatGPT checkout to partner apps in March 2026 (Section 4), so the assistant researches and the supplier transacts [17,18]. And the payment networks began building the mandate infrastructure a delegated purchase needs. Google's Agent Payments Protocol (AP2), with signed intent and cart mandates, was brought into the FIDO Alliance in May 2026 [25]; Visa's Intelligent Commerce Connect (April) is a pilot that accepts agent-initiated payments across AP2, OpenAI/Stripe's ACP and Visa's Trusted Agent Protocol through one integration, and Visa announced on 10 June that Intelligent Commerce would be integrated into OpenAI experiences with user-set spend caps, merchant-category limits and approval thresholds [23]; Mastercard's Agent Pay is in controlled pilots in Singapore and Malaysia and added an agent-to-agent variant in June [24]. The common design principle is a tokenised credential scoped to a specific agent, merchant and window, with issuer-side risk scoring. These are convergence experiments and pilots, not generally available travel-payment infrastructure; the checkout step is *becoming* a protocol choice, and a harness built today should expect to change rails. The recommendation to keep a human confirmation before each commit rests on two things: the networks' own approval-threshold designs [23], and the European trust data [3]. It does not rest on any platform having removed confirmation and reversed course; the pre-release draft implied that and was wrong.

8.4 What the benchmarks say

The academic literature on travel-planning agents has moved enough since June to change the design advice. TravelPlanner (Xie et al., ICML 2024) tests tool-grounded multi-day itinerary construction under hard constraints; pure LLM agents scored in low single digits at launch, and it is now regarded as largely solved, with reported pass rates above 97% — but by solver-assisted, neuro-symbolic methods (an SMT-based formal-verification framework) rather than by prompting alone [26]. The lesson generalises: where the task has hard constraints (budget, opening hours, travel time), the reliable harness pairs the model with a solver or verifier rather than asking it to reason its way to feasibility. The frontier has moved to what a solver cannot do. GroupTravelBench (May 2026) introduces 650 multi-user tasks in which one agent must elicit private preferences from several travellers over multiple turns and resolve conflicts between them, tracking must/prefer/avoid/reject tiers deterministically; it reports plan validity below 12% for the systems tested [27]. Its authors also summarise the prior single-user multi-turn benchmark as approaching 80% for the strongest models; that figure is their characterisation of another paper's result and is cited here as such [27]. VeriTrip (May 2026) evaluates verifiable planning over unstructured web corpora; TravelEval (June) and TREK (July) extend evaluation to richer reasoning and feasibility; TourPlanner (January) proposes a competitive-consensus multi-agent method with constraint-gated reinforcement learning; the WandaPlan work on fraud detection shows agents relaying third-party content are vulnerable to injected inducements; and ItiNera (EMNLP 2024) combines spatial optimisation with LLMs for urban itineraries [28]. Benchmark results establish relative capability, not production reliability.

Across production systems and benchmarks the pattern is the same: the model is the cheapest component. The value, and the difficulty, sit in the harness — grounding, constraints, preference elicitation, verification, and the commercial access that makes tools usable at all.

8.5 Anatomy of a production travel harness

Table 1 summarises the layers the systems above have in common, with public examples. It is a synthesis of vendor descriptions and academic work, not a disclosed architecture; items marked * are announced, wait-listed or in pilot rather than generally available.

Layer	What it does	Public examples (2026)
Orchestrator	Decomposes the request, routes to specialists, synthesises the answer; holds conversation state and memory.	Penny's coordinating layer (Claude for reasoning) [20]; Ask Maps multi-step tasks [6]; Mindtrip's planner [4].
Specialist agents	Narrow agents per vertical or task: hotel search, flights, cars, servicing, review summarisation.	Penny's 10+ agents [20]; Sabre's proxy / hotel-ops / visa / expense agent use cases [21]; GetYourGuide operator reply agent [13].
Tool & inventory layer	Live availability, prices and content reached through standard tool interfaces.	Sabre MCP server [21]; Amadeus toolkit via Claude Code* [22]; apps in ChatGPT (MCP) [17]; Viator in Gemini [15] and Claude [16].
Grounding & verification	Constraint solvers, retrieval over curated corpora, citation checks; catches infeasible or fabricated plans.	SMT-assisted TravelPlanner methods [26]; MUSEO's closed museum corpus [9]; VeriTrip-style verifiability [28].
Preference & memory	Elicits and stores traveller preferences across turns and trips; resolves group conflicts.	Penny preference learning and memory [20]; GroupTravelBench evaluates the group case [27].
Transaction & mandate	Scoped, tokenised authority for an agent to pay, with caps, categories and approval thresholds.	AP2 (FIDO)* [25]; Visa Intelligent Commerce Connect*, Trusted Agent Protocol* [23]; Mastercard Agent Pay* [24].
Guardrails & disclosure	Human confirmation before commit, injection resistance, and disclosure that the user is talking to an AI.	Visa approval thresholds [23]; WandaPlan fraud robustness [28]; EU AI Act Art. 50 from Aug 2026 [12].

Table 1. Anatomy of a production travel harness, with public examples. Author's synthesis; * = announced, wait-listed or pilot.

8.6 Implications for an in-destination harness

Every production harness described above is a booking-layer system: its specialists search inventory and its mandates authorise payment. An in-destination harness has a different job and therefore a different shape. Its tools are opening hours, event feeds, transit, weather, queue state and a curated institutional or destination corpus, not fares; its hard constraints are time, walking distance and accessibility rather than budget; its verification problem is interpretive accuracy over a trusted corpus, which is where the museum evidence and the Article 50 regime both point; and its transaction step, where it exists at all, is a ticket or a table rather than an itinerary. Four things carry over. Use a solver or verifier for feasibility rather than trusting the model [26]. Design preference elicitation for groups, because the benchmarks show this is where agents are weakest [27]. Treat the tool layer as MCP-first so that inventory partners can be added as connectors rather than integrations, while recognising that commercial access, not protocol support, is the gating factor. And keep a human confirmation at every commit, on the trust evidence [3]. Where an in-destination harness can differentiate is in the two layers the booking harnesses barely touch: grounding in a destination's own knowledge, and the mid-visit loop in which the plan is revised as the day unfolds. Neither is addressed by the incumbents' announced architectures, and the benchmarks for both (dynamic disturbances, mid-trip re-planning) are only now appearing [28].

9. Synthesis: What the Summer Changed

The June draft's central line holds. Across the in-destination layer, AI separates: discovery and engagement in heritage and destinations; the booking step in restaurants and attractions, where a platform is the active party; the marketplace in tours; back-office in car hire. The summer's indicators — a vendor-reported 17% activity-booking rate against 69% planning use, LLM-referred room nights well under 1% at the world's largest OTA, and European travellers who report using AI in destination chiefly for real-time recommendations — are consistent with the discovery-not-bypass thesis. They are a synthesis of surveys and disclosures, not a test of it. Three things are nonetheless different.

The discovery layer is now contested by platforms, not just enriched by institutions. In June the parties shaping in-destination discovery were mostly venues and a scatter of vendors. By September, Google (Ask Maps, Gemini connected apps), OpenAI (apps in ChatGPT), Anthropic (Tripadvisor and Viator connectors) and OpenTable (network-wide Concierge) are all competing to own the moment a visitor asks what to do next. The venue's content is the raw material; the contest is over whose interface delivers it.

Fulfilment stays with the supplier; discovery, data and margin may not. Direct measurement of native in-assistant booking remains limited, and most documented in-destination integrations hand off to a marketplace or platform; the museum, the operator and the restaurant remain the party that delivers. That is not the same as being unaffected. On the definition in Section 1, discovery control, customer ownership and transaction economics can shift toward the assistant or the integrated marketplace while the venue remains the supplier — the pattern already visible in restaurant reservations, where the platform decides which assistants reach the book. The intermediary most exposed is the middle-tier aggregator with neither proprietary inventory nor an assistant integration; the venue's exposure is subtler and runs through attribution and data rather than through the transaction.

DMOs are the new front, and trust is the currency. Mindtrip's announcement cadence shows the well-funded end of the market adding destinations roughly every three weeks; Google's Ask Maps shows the platform end moving into the same space; and Phocuswright's European data shows travellers using AI more while trusting it less than people. The opportunity for a destination or heritage organisation is to be the trusted source the assistants draw on; the risk is that a platform's aggregated substitute gets there first.

For an organisation building AI tools in heritage or for destinations, the implication has sharpened rather than changed. The genuine risks remain interpretive accuracy, the authority of AI-generated commentary, data governance and model-provider dependence, to which this version adds regulatory disclosure (Article 50), public reception on labour grounds (the MUSEO episode) and channel dependence. Section 8 adds the engineering corollary: the model is the cheap part, and the defensible work is in grounding, constraint handling, group preference elicitation and commercial tool access.

Cautions. Much of the new material is vendor-published: GetYourGuide, Tripadvisor, OpenTable, Priceline, Sabre and above all Mindtrip have commercial reasons to publish the figures they publish, and the Mindtrip/MMGY study is a US sample with unpublished wording. Phocuswright's European report is the most credible independent base; its full methods sit in the paid report and were not inspected, so only the publisher's public releases are relied on. Adoption evidence measures discovery far better than completed transactions. Several harness components are announced or piloted rather than available. The June draft's per-traveller spend and commission figures rested on a vendor template [43] and are not reproduced. And the best-documented players remain the large platforms; the small institutions most affected file the least public data.

10. Change Log: Corrections Made Before Release

Three rounds of independent audit were run on pre-release drafts (7 September 2026; not circulated). Statements found to be misstated, overstated or unsupported are listed with their corrections. The log is retained so that the corrections are auditable.

Pre-release draft said	Finding	Version 1.0 now says
OpenAI launched native Instant Checkout for travel in late 2025 and withdrew it by March 2026.	Misstated. Instant Checkout was retail (Etsy, Shopify); travel came via separate apps. No native travel checkout launched [17,18].	Retail Instant Checkout launched Sept 2025; travel apps Oct 2025; checkout moved to apps March 2026. (\$4, \$8.3)
The largest platform tried removing human confirmation and reversed course.	Misstated. OpenAI's launch required user confirmation of order and payment [17].	Claim removed; confirm-before-commit now rests on network approval designs and trust data. (\$8.3)
Springer museum study: nine-object grounded corpus incl. 3D models; no hallucinations.	Misstated. Three artworks; GPT-4 general knowledge; 56-participant study; expert found factual soundness with terminology gaps [11].	Accurate description; explicitly does not validate a grounded production guide. (\$3.2)
Amadeus 'moved its content into Claude Code'.	Overstated. Announced collaboration; toolkit wait-listed, plugin under development [22].	Described as announced, not available; marked * in Table 1. (\$8.2)
MUSEO launched in August; first well-documented labour backlash.	Misdated and not auditable. Pilot public by early 2026; vendor note 10 July; reaction August [9,10]. Firstness not established.	Dated correctly; priority claim removed. (\$3.1)
Car hire: 'no material change'.	Not auditable as a market-wide negative.	Restated as 'no new evidence found' with the search scope disclosed. (\$7)
Phocuswright report 'published July'.	Publisher listing dates it June 2026 [1].	June, with AI-specific releases in September. (\$2.3)
LTM guide 'an order of magnitude below a traditional refresh'.	Unsourced comparison.	Removed; cost stated as the museum's own figure. (\$3.1)
OpenTable APIs 'partner-gated with no self-serve access'.	Developer portal exists; production access is gated [29].	'No open public API'. (\$5)
Figures 1–3 as bar charts.	Implied measurement or common denominators where none exist.	Fig. 1 qualitative matrix; Fig. 2 source cards; Fig. 3 two labelled snapshots.
Grouped prose references; June draft sources cited only as a block.	Not reproducible at claim level.	Numbered claim-level citations with evidence-type labels; June draft sources itemised, six upgraded to primary records, remaining aggregators flagged. (\$11)
Round 2. OpenTable 61%: 'methodology unpublished'; release dated 15 July.	Misstated. Methodology is published (Ripple Research, n=1,500 hotel-restaurant diners, 28 May–2 Jun 2026); release dated 14 July.	Methodology stated; sample limitation noted; dates corrected. (\$5, item 30)
Figure 2 titled 'adoption-to-transaction gap'.	Overstated: the three measures observe no common funnel.	Retitled 'three non-comparable indicators'; caption says they do not establish a conversion gap. (\$4)
'The transaction gap protects venues'; fulfilment treated as evidence of non-bypass.	Definition of disintermediation too loose.	Five-part definition added (\$1); synthesis rewritten: fulfilment stays with supplier, discovery/data/margin may not. (\$9)
'First sign'; 'where hallucination risk is highest'; 'a deal a fortnight'.	Unsupported superlatives; cadence miscounted (seven announcements, ~3 weeks apart, none in July).	'A sign'; 'risk may be higher because model knowledge is thinner'; 'roughly every three weeks; announcement cadence is not contract cadence'.
Item 35 omitted third author and used page range 1–25; item 36 'Predescu et al.'.	Bibliographic errors.	Trichopoulos, Ordoumpozanis & Caridakis, 14:1–14:25; Derda & Predescu with DOI.

Pre-release draft said	Finding	Version 1.0 now says
Items 31 and 35–43 listed but not cited in text; no hyperlinks; no bookmarks.	Reproducibility.	All items now cited at specific claims; DOI/URL links added where verified; PDF bookmarks added.
Round 3. In-assistant transactions 'appear to remain rare'; 'the assistant rarely completes the booking'.	Frequency claim not measured by any cited indicator (see Figure 2 caption).	'Direct measurement of native in-assistant booking remains limited; most documented integrations hand off to a marketplace; Google's AI Mode hotel booking is a counterexample.' (\$4, \$9)
'The most actively contested segment'; 'the best-funded vendor in the segment'.	Unsupported rankings.	'One of the most visibly contested segments in the evidence reviewed'; 'a well-funded vendor'. (\$1 note, \$2.1)

11. Sources

Each source is labelled by evidence type: **[Peer-reviewed]**, **[Preprint]**, **[Law/guidance]**, **[Earnings]** (investor disclosure), **[Product record]** (official product announcement), **[Vendor release]** (company press release or blog), **[Vendor survey]** (company-commissioned research), **[Trade]** (trade press), **[Press]** (general news). Vendor material verifies what a company said or launched, not adoption, accuracy or return. Verified against public records to 7 September 2026. Links are given where a stable URL or DOI was verified during drafting; where none appears, the item was seen via search results or trade coverage and should be located by title. Items 32–43 are the June draft's sources, itemised from its reference list. Six that rested on secondary or tertiary sources (items 35, 36, 37, 39, 41 and 42) have been re-verified and replaced with primary records, marked ✓; the remainder were not re-verified and several are secondary aggregators, flagged as such.

- [1] Phocuswright Research, *Europe Consumer Travel Report 2026* (UK, France, Germany; fielded April 2026; published June 2026). Paid report, not inspected; cited only through the publisher's public releases (items 2 and 3). [Independent research]
- [2] PhocusWire, 'AI isn't just for research anymore: European travelers say it's shaping trip decisions', 1 Sep 2026, presenting item 1. [Trade] <https://www.phocuswire.com/news/technology/ai-shaping-trip-decisions-phocuswright-research-2026>
- [3] Phocuswright press release, 'European travelers are using AI more, but trusting it less', 4 Sep 2026. [Publisher release]
- [4] Mindtrip press releases: Europe launch (Madeira Promotion Bureau, Norwegian Travel Cluster), 3 Jun 2026; Hilton Head, 10 Jun; Savannah/Tybee Island, 16 Jun; Meet Boston and WhatsApp integration, 30 Jun; San Diego Tourism Authority, 28 Aug; Meet Minneapolis, 2 Sep; San Francisco Travel, 29 Apr 2026. [Vendor release]
- [5] Mindtrip with Sabre, Sojern and MMGY Travel Intelligence, AI readiness study; fielded April 2026 (n=1,006 US adults; n=103 DMO leaders); release 20 Aug 2026 via Hotel-Online. Question wording unpublished. [Vendor survey] <https://www.hotel-online.com/news/new-research-reveals-travelers-are-two-years-ahead-of-the-travel-industry-on-ai>
- [6] Google, Ask Maps update (hotel, event and dining discovery; 150+ countries; agentic features US-only), 6 Aug 2026; PhocusWire coverage same date. [Product record; Trade]
- [7] Google, agentic hotel booking in AI Mode (US English), 27 Aug 2026; Skift coverage same date. [Product record; Trade]
- [8] MuseumNext, 'Bringing a museum collection to life with conversational AI' (London Transport Museum AI Digital Poster; Museum AI Summit 2026), Jun 2026. [Trade case study; cost is the museum's own figure]
- [9] Ask Mona, 'An AI to empower the visitors at the Canadian Museum of History' (MUSEO), 10 Jul 2026 [Vendor case note] <https://www.askmona.ai/blog/museo-canadian-museum-history-conversational-ai>; The Globe and Mail, 'Inside the Canadian Museum of History's experiment with artificial intelligence', early 2026 [Press] <https://www.theglobeandmail.com/business/article-inside-the-canadian-museum-of-historys-experiment-with-artificial/>
- [10] Postmedia (Ottawa Citizen), 'I tried the Canadian Museum of History's new AI chatbot', Aug 2026 <https://ca.news.yahoo.com/tried-canadian-museum-historys-ai-080044290.html>; CBC News video, Aug 2026. [Press]
- [11] Dibitonto M., Ferrato A., Limongelli C., Patroni O.C., 'Museum audio guides generation using visitor categories and large language models', *Multimedia Systems* (2026), Three artworks; GPT-4; n=56; expert validation. [Peer-reviewed] [doi:10.1007/s00530-026-02494-5](https://doi.org/10.1007/s00530-026-02494-5)
- [12] Regulation (EU) 2026/1744 (Digital Omnibus on AI), EUR-Lex, in force 27 Jul 2026; Regulation (EU) 2024/1689 (AI Act) Arts. 50 and 99; European Commission guidance on Article 50 transparency obligations (2026). [Law/guidance]
- [13] GetYourGuide press release, Spring 2026 product release, 22 Apr 2026, previewing research with Arival (May 2026); Skift coverage 22 Apr 2026. Consumer sample design undisclosed. [Vendor release; Vendor survey]
- [14] GetYourGuide, *Spring 2026 Travel Experience Trend Tracker* (n=505 operators), May 2026. [Vendor survey]
- [15] Tripadvisor MediaRoom: 'Building what's next in experiences: summer update', 10 Jun 2026; 'Viator becomes Google's first connected app for travel experiences' (425,000+ experiences), 30 Jul 2026. [Vendor release]
- [16] PhocusWire, 'Tripadvisor, Viator expand AI partnerships with Claude, Alexa+', 23 Apr 2026; Anthropic connector directory listing Viator (experiences) and Tripadvisor (hotels), checked 7 Sep 2026. [Trade; Product record]
- [17] OpenAI, 'Buy it in ChatGPT: Instant Checkout and the Agentic Commerce Protocol', 29 Sep 2025 (retail; user confirms order and payment); OpenAI, 'Introducing apps in ChatGPT and the Apps SDK' (built on MCP), 6 Oct 2025. [Product record]
- [18] Skift, 'ChatGPT bails on transactions — good news for Expedia and Booking', 5 Mar 2026, quoting an OpenAI spokesperson <https://skift.com/2026/03/05/openai-chatgpt-checkout-walkback/>; PhocusWire, 'OpenAI's shift shows travel is too complex for quick-fix distribution', 10 Mar 2026 <https://www.phocuswire.com/news/technology/openai-chatgpt-instant-checkout-travel-intermediaries>. [Trade]
- [19] Booking Holdings, Q2 2026 results and earnings call, 4 Aug 2026 (LLM-referred room nights 'significantly less than 1%'; attraction tickets double-digit growth; Connected Trip low double digits). [Earnings]

- [20] Priceline press release, 'Priceline's Penny goes fully agentic', 3 Jun 2026; Skift, 'Priceline gives Penny a multi-agent makeover', 3 Jun 2026. Performance claims without effect sizes. [Vendor release; Trade]
- [21] Sabre press releases: agentic APIs and MCP server, Sep 2025; BizTrip AI partnership, 14 Jan 2026; Linex Travel production deployment (12 countries), Jun 2026. [Vendor release]
- [22] Amadeus blog post by L. Viguié on collaboration with Anthropic, 3 Sep 2026; Skift, 'Amadeus brings travel tech into Claude', 3 Sep 2026; Business Travel News, 3 Sep 2026. Toolkit wait-listed; Cowork plugin under development; terms unpublished. [Vendor release; Trade]
- [23] Visa, Intelligent Commerce Connect (pilot), 9 Apr 2026; Visa and OpenAI partnership announcement, Visa Payments Forum, 10 Jun 2026. [Vendor release]
- [24] Mastercard, Agent Pay controlled pilots (Singapore, Malaysia) and Agent Pay for Machines, Jun 2026. [Vendor release]
- [25] FIDO Alliance, trusted AI agent standards work incorporating AP2, May 2026. [Standards body]
- [26] Xie J. et al., 'TravelPlanner: a benchmark for real-world planning with language agents', ICML 2024 [Peer-reviewed]; Hao Y. et al., formal-verification (SMT-assisted) framework for TravelPlanner reporting ~97% pass rate, 2025 [Preprint].
- [27] Cheng X., Hu Y., Zheng L., Pan Z., Li X., Liu Y., 'GroupTravelBench: benchmarking LLM agents on multi-person travel planning', arXiv:2605.25200, May 2026 (650 tasks; plan validity <12%; summarises prior multi-turn benchmark at ~80%). [Preprint] <https://arxiv.org/abs/2605.25200>
- [28] VeriTrip, [arXiv:2605.28683](https://arxiv.org/abs/2605.28683) (May 2026); TravelEval, [arXiv:2606.01046](https://arxiv.org/abs/2606.01046) (Jun 2026); TREK, [arXiv:2607.26977](https://arxiv.org/abs/2607.26977) (Jul 2026); Wang Y. et al., TourPlanner, [arXiv:2601.04698](https://arxiv.org/abs/2601.04698) (Jan 2026); WandaPlan / 'Is your LLM-based multi-agent a reliable real-world planner? Fraud detection in travel planning', [arXiv:2505.16557](https://arxiv.org/abs/2505.16557) (2025) [Preprints]; Tang Y. et al., ItiNera, EMNLP 2024 Industry Track, [doi:10.18653/v1/2024.emnlp-industry.104](https://doi.org/10.18653/v1/2024.emnlp-industry.104) [Peer-reviewed].
- [29] OpenTable press release, 'Book the Restaurant, Then the Room: OpenTable Reveals the 2026 Top 100 Hotel Restaurants in America', 14 Jul 2026, including the Concierge homepage announcement (Concierge originally launched on restaurant profiles in 2025) <https://www.opentable.com/blog/press/page/hotel-restaurants-2026/>; OpenTable developer portal FAQ and API partner terms (production access application- and partner-gated). [Vendor release; Official documentation]
- [30] OpenTable consumer research released with item 29, 14 Jul 2026: online survey by Ripple Research of 1,500 US residents who had dined in a hotel restaurant within the past five years, major cities weighted at 200, fielded 28 May–2 Jun 2026, MRS/ESOMAR-compliant; 61% chose a destination for its food scene; 51% plan to use AI tools to discover and book restaurants while travelling. Sample is hotel-restaurant diners, not the general population. [Vendor survey, methodology published] <https://www.opentable.com/blog/press/page/hotel-restaurants-2026/>
- [31] PhocusWire, 'AI transformation in travel: Piero Sierra of Skyscanner', Sep 2026 (Explore with AI beta; DMO content in AI results). [Trade]
- [32] **June draft sources, itemised.** The items below are transcribed from the June draft's reference list and labelled by evidence type; six have been upgraded to primary records (marked ✓). Where an aggregator remains, the underlying primary source has not been retrieved. MuseumNext, 'How Are Museums Using Artificial Intelligence?' (Louvre 'Leonardo' assistant; British Museum AR app and chatbot), 2024–2025. [Trade]
- [33] Museumfy, Metropolitan Museum of Art chatbot drawing on 45,000+ artworks, 2024–2025. [Vendor blog; secondary]
- [34] Mapsted and Travel And Tour World, Museum of Tomorrow (Rio de Janeiro) IRIS+ chatbot accessibility upgrade (sign-language translation, audio description), 2024–2025. [Vendor blog; Trade]
- [35] ✓ Trichopoulos G., Ordoumpozanis K., Caridakis G., 'An Evaluation of LLM-based Chatbots for Enhancing the Visitor's User Experience at Cultural Exhibits', *ACM Journal on Computing and Cultural Heritage* 19(1), Art. 14, 14:1–14:25, published 22 Jan 2026, [doi:10.1145/3775062](https://doi.org/10.1145/3775062). Three field tests (digital art exhibition, painting gallery, Archaeological Museum of Mytilene); mixed methods; LLM guides outscored a static chatbot on engagement and meaningful experience, with context-awareness and hallucination named as unresolved. [Peer-reviewed]
- [36] ✓ Derda I., Predescu D., 'Towards human-centric AI in museums: practitioners' perspectives and technology acceptance of visitor-centered AI for value (co-)creation', *Museum Management and Curatorship* 40(4), 532–554, 2025 (open access, CC BY-NC-ND), [doi:10.1080/09647775.2025.2467703](https://doi.org/10.1080/09647775.2025.2467703). Proposes the Human-Centric AI Governance Model for Museums. [Peer-reviewed]
- [37] ✓ Mariani M.M., Hashemi N., Wirtz J., 'Artificial intelligence empowered conversational agents: a systematic literature review and research agenda', *Journal of Business Research* 161 (2023), 113838 (open access), [doi:10.1016/j.ibusres.2023.113838](https://doi.org/10.1016/j.ibusres.2023.113838). Identifies consumer trust in conversational agents as a dominant research theme. [Peer-reviewed]
- [38] AITourPilot and Visit Widget ('Via' destination chatbot) product material, 2025. Engagement percentages treated as promotional and not relied upon. [Vendor]
- [39] Dittofi, 'Viator competitors' (GetYourGuide, Klook, TUI Musement; note on pen-and-paper operators), 2025 [Secondary aggregator]; GetYourGuide ChatGPT plugin listing (75,000+ activities), getyourguide.com [Vendor]; ✓ GetYourGuide

press blog, milestone post by J. Reck, 21 Oct 2025 (record Q3 2025: 10 million experiences booked, GMV +30% year on year; adjusted-EBITDA profitable; revenue approaching €1 billion) getyourguide.press; Skift coverage same date skift.com [Vendor release; Trade]. Replaces the Wikipedia citation in the June draft.

- [40] AlInvest, 'How OpenTable's AI Concierge is Reshaping Restaurant Tech' and 'Google's AI-Driven Restaurant Reservation Ecosystem' (Copilot/Alexa+ integrations; ~46% US share; Yelp reservation growth), 2025 [Secondary aggregator]; Quiver Quantitative, SoundHound–OpenTable in-vehicle agent, Dec 2025 [Secondary]; Booking Holdings Form 10-K 2025 (OpenTable), filed 2026 [Earnings/filing]; ✓ OpenTable official plans page, restaurant.opentable.com/products (Basic US\$149/month plus US\$1.50 per network cover; Core US\$299 plus US\$1.00; Pro US\$499 plus US\$1.00; 2% service fee on prepaid experiences and ticketing; US pricing, checked Sep 2026) <https://restaurant.opentable.com/products/>. Replaces the trybytes.ai citation in the June draft. [Official pricing page]
- [41] ✓ Booking Holdings, Q4 and full-year 2025 earnings call and prepared remarks, 18 Feb 2026 (attraction tickets booked grew nearly 80% year on year from a relatively small base; Connected Trip transactions grew in the high-20% range and are a low-double-digit share of Booking.com transactions; 68 million airline tickets, +37%). Transcript: fool.com. Replaces the Intellectia secondary analysis in the June draft. [Earnings]. Dittofi on Klook theme-park and event partnerships, 2025 [Secondary aggregator].
- [42] Kuchoriya TechSoft, AI in car-rental operations and dynamic pricing, 2025 [Vendor blog]; StockCounterparts and Booking Holdings Form 10-K on rental cars within Connected Trip, 2025–2026 [Secondary; Earnings/filing].
- [43] Conferbot, local-attraction hotel-chatbot template (per-guest per-day discretionary spend US\$85–140; 15–25% tour/experience commissions), 2026. Basis of the June draft's Figure 2. [Vendor template; illustrative ranges, not audited — the weakest source in either draft and the reason that figure is not reproduced here.]

Version 1.0 (first release), September 2026. Companion piece to 'Under Threat? AI Chatbots, Travel Tools, and the Uneven Exposure of Travel-Industry Segments' (unreleased draft, June 2026). Descriptive in intent; interpretation is the author's unless attributed. Drafting history: June 2026 (internal draft); 7 September 2026 (internal update, audited); version 1.0 released following audit corrections. Legal references are informational analysis, not legal advice.