

Beyond the Booking Funnel

Five findings on AI in the in-destination layer, from a September 2026 evidence review

One-page summary · Richard Hensler, Mariners Compass Ltd · September 2026 · summarises the full paper, v1.0, 17 pp.

The booking layer asks who captures the transaction. The layer that starts when the visitor arrives (museums, tours, restaurants, attractions and the destination itself) asks who shapes discovery and the day. Most discussion of AI in travel is about the first. The evidence about the second is less comfortable.

1 The 'under 1%' figure measures referrals, not influence

Booking Holdings' disclosure that LLM-referred traffic is well under 1% of room nights (Q2 2026) is widely read as reassurance. In the same summer, Phocuswright's public release reported that among surveyed European travellers who use AI, 88% in Germany and 92% in France say it shaped an actual trip decision. The two figures have different populations, denominators and outcomes, and they are not a funnel. Together they suggest that referral attribution captures only part of AI's influence. On the paper's definition, a venue can keep all of its fulfilment while discovery control, customer data, attribution and margin move toward the assistant. None of those four appears in the referral metric.

2 The documented harnesses stop at the hotel door

The publicly documented harnesses the paper reviews (Priceline's ten-plus agents, Sabre's MCP deployment, the ChatGPT and Gemini apps) are booking-layer systems: specialists that search inventory, mandates that authorise payment. None is described as re-planning the day as it unfolds, using hours, queues, weather and a curated local corpus as its tools. Yet Phocuswright's public release finds that this is where travellers use AI in destination. Real-time recommendations lead in every market surveyed (53% of AI-using UK travellers), with trip management twenty points behind.

3 The protocol is standardising; access is the moat

Sabre (in production), Amadeus (announced, wait-listed) and OpenAI's apps have converged on the Model Context Protocol as the way agents reach inventory. The protocol is not the hard part. Commercial access is, and it is gated everywhere: Amadeus's terms are unpublished, Sabre's server sits inside enterprise contracts, and OpenTable has no open public API. A destination or museum reaches the visitor's assistant through whoever holds the keys, unless its own knowledge becomes the source the assistant is grounded in.

4 Trust is moving the other way, and the trusted asset is sitting unused

Phocuswright's public release reports that Europeans are using AI more and trusting it less than friends, family and human advisers. The paper infers from this (an inference, not a survey finding) that the asset travellers still trust is an authoritative local source, which is what a destination or heritage institution is. Mindtrip's commissioned survey of 103 US DMO leaders found that 22% offer an AI trip planner, while 63% of surveyed travellers would use one. Being the source matters more than owning the AI.

5 Your visitor chatbot has had a legal disclosure duty since 2 August

The Digital Omnibus (Regulation 2026/1744) deferred the AI Act's high-risk rules to 2027 and 2028 but left Article 50 on its original date. Since 2 August 2026 the provider of an EU-facing chatbot must make clear that the user is dealing with an AI, unless that is already obvious to a reasonably informed person. Generative output needs machine-readable marking, with a grace period for existing systems to 2 December. Fines reach €15m or 3% of turnover. The duty is easy to meet, but a museum or destination using a vendor's system should have the contract say who the provider is and who carries it.

Questions to ask your AI vendor	
What is it grounded in: our validated content, or the model's general knowledge? (Finding 4)	Which assistants can reach our content, and who holds the keys? (Finding 3)
Does it keep working after arrival, and what does it do when the day changes? (Finding 2)	Where does a person confirm before anything is committed? (Finding 4)
Who is the Article 50 provider, and what visitor data do we own? (Finding 5)	What will we be able to measure: usage, and whether AI-shaped decisions show up in our attribution? (Finding 1)
Where this leads. PersonalTourister is built to answer those questions. It is an AI concierge that answers by retrieval from a destination's own curated content, so the destination's knowledge is the source and the model only composes the reply. It profiles the visitor in conversation and is designed for the hours after arrival. With Giorgio, who leads our Italian destination work, we are piloting it with municipalities in northern Italy, which keep their content and their visitor data. We would like to run a comparable pilot with a UK or Irish destination.	

Phocuswright figures are from its public releases of 1 and 4 September 2026, not the paid report. Mindtrip, GetYourGuide and Priceline figures are vendor-published. Full citations are in the paper (17 pp.; three audit rounds by an independent AI reviewer).