

HYB26

Technology Edition

AI Everywhere



The Hotel Yearbook

Foresight and innovation in the global hotel industry

HY8



From Search to Synthesis: Visibility in an Answer- Based Internet

Pre-Stay

Alessio Re
Co-founder, Elegia

Σ Elegia

The pre-stay booking funnel has fundamentally shifted, and most hotel tech stacks are nowhere near ready for it. Alessio Re maps the three-layer challenge hotels now face - getting cited by AI systems, owning accurate and structured data, and having a vendor stack capable of converting agent-driven discovery into actual bookings - and argues that treating this as an SEO problem with new vocabulary will not get the industry far enough, fast enough.

The funnel for finding and booking a hotel has split across two surfaces, and most hotel tech stacks are not ready for either. In March 2026, Walmart's EVP for design and product, Daniel Danker, told WIRED that purchases inside ChatGPT through OpenAI's Instant Checkout had converted at one-third the rate of click-outs to [Walmart.com](https://www.walmart.com) on the same catalogue (Dave 2026). Walmart ran 200,000 SKUs through Instant Checkout from November until OpenAI pulled it back, pivoting to merchant-controlled apps inside ChatGPT, the architecture Walmart's Sparky and Accor's ALL app already use. Hotels are not selling USB cables. They sell rate-and-availability bundles that change hourly, with cancellation policies, loyalty tiers, package inclusions and city fees that a chatbot flattens or gets wrong. If a 200,000 SKU retailer failed at in-chat conversion, hotels with messy data and a locked vendor stack will fail worse.

Pre-stay visibility in 2026 rests on three things:

- Content architecture that earns a citation in an AI answer
- A data layer that lets a hotel describe itself accurately to a model
- A vendor stack that lets an agent act on what it finds

Too many hotels still treat this as an SEO problem with new vocabulary, which solves only a fraction of the first item and none of the other two. A property that does all three has a chance of getting cited by the models that matter and converting that interest into a direct booking.

THE CITATION LAYER HAS DECOUPLED FROM THE RANKING LAYER.

January 2026's upgrade of AI Overviews to Gemini 3 changed how citations get selected. Ahrefs measured the share of AI Overview citations that also ranked in the top 10 organic results, dropping from 76% in July 2025 to 38% in February 2026, across 863,000 keywords and 4 million URLs; BrightEdge put the overlap closer to 17% (Southern 2026). The numbers diverge, but they point in the same direction: ranking well is no longer a proxy for being cited.

The main mechanism is called query fan-out. Gemini, the AI system, breaks the user's request into smaller questions called sub-queries. It then retrieves information for each part and cites the web pages that consistently answer these sub-questions, not just the page ranked first for the original query. This means Overviews are built using sources that cover specific parts of the question, such as a TripAdvisor thread or a YouTube creator's video.

The citation pool itself is not where most hotels expect to find themselves. The 5W AI Platform Citation Source Index synthesized 680 million citations across ChatGPT, AI Overviews, Perplexity, Gemini, and Claude (5W 2026): Reddit accounts for 40% of LLM citations, YouTube is the most-cited domain in Google's AI surfaces, and the top 15 domains capture 68% of the pipeline.

A hotel can earn a citation directly by publishing accurate, dated content answering a granular question: pet policy with the fee and breed restriction, parking with the height clearance, late check-in with the cut-off, and contact method. A page in plain FAQ format with structured data markup reads the same way to a human and to a language model.

Google's May 2026 AI Overviews redesign added five citation surfaces, including hover previews, a subscribed-publisher label, and a Community Perspectives panel quoting Reddit and forum posts with creator handles attached (Budaraju 2026). The shape of what gets cited shifts month to month; anything in this section is calibrated to mid-2026.

DATA OWNERSHIP IS THE PILLAR MOST HOTELS SKIP.

A model can only cite what it reads accurately. A hotel can only be read accurately if it owns the canonical version of its information. For most hotels, that version does not exist. Pricing lives in the revenue management system. Availability is in the channel manager. Descriptions appear in the booking engine or PMS. Policies are in a Word document. Loyalty terms reside on a white-label third-party platform. None of those systems talks to each other with full fidelity. The OTA listing is often cited because it is structured and updated against a known schema.

The 2026 HotelRank study audited 121,000 hotel homepages across seven countries and found 36% with no structured data, and 10% with what it classified as good Hotel or LodgingBusiness schema (HotelRank 2026). The fix is rarely a technology problem; it is a question of who owns and updates the data. Independent hotels and smaller groups need a warehouse-first architecture in which the PMS, booking engine, email service provider, and consent records flow into a layer they control. Marriott earmarked \$1.1 billion for 2026, with 35 to 40% going towards replatforming its property management, reservations, and loyalty systems, while the CEO described the company as "pulling into the players' parking lot" on agentic AI (Marriott International 2026). Smaller operators will never spend that, but the architecture is achievable on a smaller budget.

European regulation is moving in two directions, and neither addresses the underlying problem. On 16 April 2026 the Commission's preliminary findings against Google, part of a specification proceeding opened on 27 January, proposed Google share ranking, query, click and view data with third parties on fair, reasonable and non-discriminatory (FRAND) terms, with AI chatbots that have search functionality named "data beneficiaries"; a binding decision is due by 27 July 2026 (European Commission 2026). The intent is to open AI search, but forcing those signals to flow to more third parties widens the privacy surface, which the proceeding does not address.

The AI Act's Article 50 transparency obligations, from 2 August 2026, require hotels to disclose AI-generated content and machine-mark synthetic imagery. Both regulations assume the hotel knows what its data is, where it lives, and who has it.

THE VENDOR STACK DETERMINES WHETHER VISIBILITY CONVERTS.

Discovery is one half of the pre-stay funnel. The other half is conversion: whether an agent who has identified the right hotel can complete the booking without the flow falling over. This is where most hotels lose the agent traffic they've earned, because the vendor stack governs whether the booking completes.

Two early-stage technical standards saw progress in February 2026. WebMCP (a protocol that allows websites to communicate structured actions, such as searching for availability or applying rates, to AI agents, developed by Google and Microsoft) launched as a preview in Chrome 146 on 10 February. Rather than forcing AI to read website screens as humans do, WebMCP tells the agent exactly what functions are available (Bokan et al. 2026).

Cloudflare's Markdown for Agents, released two days later, provides simplified versions of webpages specifically designed for AI to read, reducing the work for AI agents by around 80% (Martinho and Allen 2026). MCP, a separate server-side protocol, connects AI assistants directly to hotel backend systems. Each vendor should prepare to adopt these protocols.

The first significant hospitality move on this architecture is Accor's. The group launched ALL inside ChatGPT on 29 January 2026, in over twenty languages, across brands from Sofitel to ibis. The user searches and configures in chat, sees rates and amenities, then, upon booking, is redirected to the ALL platform to complete the reservation on Accor's surface (Accor 2026). The handoff matters because Walmart's in-chat conversion collapsed: the chat window strips out brand context, reviews, bundle logic, upsell, and the cancellation policy travelers rely on at commitment. Chat is fine for discovery, but the actual conversion needs to land somewhere the merchant controls, and most merchants in our field have not built that yet.

Most hotels do not own their checkout flow. It sits inside a white-label booking engine widget, governed by that vendor's release cadence. The same vendor lock-in that forces hotels to route pricing through third-party feeds to metasearch and OTAs now determines whether they appear in agentic commerce. The chokepoint that prevented Book on Google from working will determine whether a hotel can support WebMCP, expose a merchant app in ChatGPT or Gemini, and route agent discovery back to a direct booking. Three questions for every PMS, CRS, channel manager, and booking engine vendor this year:

- What is your roadmap for MCP and WebMCP, and how quickly do you ship support when OpenAI, Google, or Anthropic publishes the next protocol update?
- Can your booking process be tagged so an AI agent can access it directly?

- Do I own my data, with the right to export it cleanly and without high export fees?

Vague answers mean the vendor will be the bottleneck on everything upstream.

THE NEXT TWELVE MONTHS

A few guesses for the next twelve months. I expect to be wrong about at least one.

WebMCP will likely shift from preview to general availability in Chrome and Edge by late 2026. Some booking engine vendors will support it, but most won't. Supportive vendors will become more attractive to attentive hotels.

More major chains will launch ChatGPT apps on the Accor template, and at least one large OTA will do the same. Independents and smaller groups will sit this out: their vendor stacks will not let them in.

The citation pool will continue to move away from the top 10 organic results. Hotels that publish dated, granular content on their own pages should catch some of that traffic. The ones still relying on their OTA listings as the canonical version will lose ground in places they cannot easily see.

Vendors in our field have long been slow and conservative; last year highlighted this further. This likely won't change soon, though I hope I am wrong.

REFERENCES

- Accor (2026), "Accor leads hospitality innovation with the launch of its ALL Accor app in ChatGPT", press release, available at: <https://press.accor.com/accor-leads-hospitality-innovation-with-the-launch-of-its-all-accor-app-in-chatgpt/?lang=eng> (accessed 18 May 2026).
- Bokan, D., Sagar, K., Van Opstal, H., Walderman, B., Lee, L., and Nolan, A. (2026), "WebMCP is available for early preview", Chrome for Developers Blog, available at: <https://developer.chrome.com/blog/webmcp-epp> (accessed 18 May 2026).
- Budaraju, H. (2026), "5 new ways to explore the web with generative AI in Search", Google Blog, available at: <https://blog.google/products-and-platforms/products/search/explore-web-generative-ai-search/> (accessed 18 May 2026).
- Dave, P. (2026), "Why Walmart and OpenAI Are Shaking Up Their Agentic Shopping Deal", WIRED, available at: <https://www.wired.com/story/ai-lab-walmart-openai-shaking-up-agentic-shopping-deal/> (accessed 18 May 2026).
- European Commission (2026), "Commission proposes measures to Google on sharing search engine data with third parties under Digital Markets Act", press release IP/26/825, available at: https://ec.europa.eu/commission/presscorner/detail/en/ip_26_825 (accessed 18 May 2026).

HotelRank (2026), "Hotel [Schema.org](https://hotelrank.ai/research/hotel-schema-adoption-study-2026) Adoption Study 2026", available at: <https://hotelrank.ai/research/hotel-schema-adoption-study-2026> (accessed 18 May 2026).

Marriott International (2026), "Fourth Quarter 2025 Earnings Conference Call Transcript", available at: <https://marriott.gcs-web.com/static-files/30de714f-753d-4ee1-8ea8-d55cf9d7424a> (accessed 18 May 2026).

Martinho, C. and Allen, W. (2026), "Introducing Markdown for Agents: making the web a first-class destination for AI", Cloudflare Blog, available at: <https://blog.cloudflare.com/markdown-for-agents/> (accessed 18 May 2026).

Southern, M.G. (2026), "Google AI Overview citations from top-ranking pages drop sharply", Search Engine Journal, available at: <https://www.searchenginejournal.com/google-ai-overview-citations-from-top-ranking-pages-drop-sharply/568637/> (accessed 18 May 2026).

5W (2026), "AI Platform Citation Source Index 2026", available at: <https://everything-pr.com/ai-platform-citation-source-index-2026/> (accessed 18 May 2026).

