A Milestone Research Report

Milestone

Q3 2026 edition

Hospitality AI Visibility Benchmark 2026

How ChatGPT, Gemini, Perplexity, Claude and Google AI Overviews find, name and cite hospitality businesses. Built from real AI answers.

1.38M
AI answers analyzed

1.6K
Hospitality locations measured

5
AI engines measured

In partnership with

AI Hospitality Alliance

EDNA

AVEN HOSPITALITY

Published November 2026

Contents

	Executive Summary The benchmark and key findings at a glance	03
	The Shift to AI-led Travel Planning Phocuswright, Adobe and Pew research on traveler adoption and click behavior	04
01	**Being Known by AI vs. Being Named or Cited by AI** Branded vs unbranded visibility	05
02	**Visibility Varies by Prompt Intent** How likely AI is to name or cite a property, by the kind of question asked	06
03	**Different AI Engines Trust Different Sources** Source mix by engine, hotels and resorts	07
04	**Booking and Review Sites Shape the AI Answer** The public domains engines cite for hotels	09
05	**Citing OTA Listings and Reviews Coincides with Higher Visibility** Page types linked to visibility, owned-site citations and share of voice	10
06	**Perplexity Named Hotels More Often This Quarter, Google AI Overviews Less Often** Quarter-on-quarter change by engine, July to September 2026	11
07	**No Clear Brand-Affiliation Advantage Was Established in This Sample** AI Visibility by Hotel Segment	13
	Recommendations Five priorities for hospitality marketers	14
	Methodology and Sources Sample, engines, metrics, statistics and sources	15
	About the Partners Milestone Inc., AIHA, HEDNA and AVEN Hospitality	17

Joint foreword

Milestone Inc. built this benchmark from 1.38 million live AI answers across 1,600 Milestone hospitality customer locations. The AI Hospitality Alliance (AIHA), HEDNA and AVEN Hospitality join Milestone Inc. in releasing it to the industry, giving hoteliers a shared, evidence-based view of how often AI engines name or cite properties, which sources each engine trusts, and how that is shifting quarter to quarter.

Executive summary

Milestone Research

Key Findings

A benchmark of observed AI visibility across 1,600 Milestone hospitality customer locations and five AI engines, July–September 2026. Figures describe this customer base, not the entire hospitality market. Hospitality covers 1,128 hotels and resorts and 472 RV parks, campgrounds and communities. RV parks are 30% of locations but 54% of answers. Apart from the 63% and 55% benchmarks, the findings below are for hotels and resorts only.

63%

AI Visibility (mention or citation) across all prompts

The hospitality benchmark

Across 1.38M AI answers, median location-level AI Visibility (named or cited) was 63%. The middle half of properties range from 42% to 79%.

55%

named or cited in *"best hotel..."* answers

AI names properties on branded prompts far more reliably than on unbranded prompts

On branded prompts, AI names the property in essentially every answer. On unbranded prompts, where engines choose among competing properties, the median falls to 55%.

Hotels and resorts only: median 54%.

8%

of AI citations come from the hotel's own website

Owned websites are cited often but form a small share of sources

A hotel's own site appears in 45% of its answers, yet owned sites are only 8% of sources cited. Booking sites (21%), travel media and other businesses (14% each) and review sites (11%) make up most of the rest.

49%

of unbranded hotel answers name or cite an OTA

Online travel agencies are prominent in AI hotel answers

For the median hotel, 49% of unbranded prompt answers name or cite an OTA. Booking and review sites together account for 32% of hotel citations; booking sites alone, 21%.

15%

of Google AI Overviews citations go to the hotel's own site

Google AI Overviews cites property websites more than other engines do

The highest of any engine, about 1.4 to 2.5 times the chat engines (6–11%), with the fewest OTA citations (7%).

47%

of unbranded answers name brand-affiliated hotels (55% independents)

No clear brand-affiliation advantage

Brand-affiliated hotels were named in a median 47% of unbranded prompt answers, independents in 55%. The ranges overlap: the data did not establish a reliable difference between segments.

70%

Perplexity visibility for hotels and resorts, up from 65%

Perplexity now names or cites hotels more often

Hotel visibility rose from 65% in July to 70% in September, and 72% of hotel brands saw the rise.

The shift

Travelers Now Consult AI Before They Book

Hotel discovery is moving from search results to synthesized answers. A traveler planning a weekend in Santa Fe once compared a page of links. Increasingly, that traveler asks ChatGPT or reads the AI summary above Google's results, and leaves with a short list of three names. Independent research indicates this already shapes how a significant share of travelers plan.

56%
of active US travelers used AI for at least one trip in the past year
Phocuswright, Mar 2026

5×
since 2024: 33% of travelers now use generative AI platforms for trip research
Phocuswright, Jul 2026

+194%
AI-referred visits to US travel sites in May 2026, year-over-year
Adobe, Jun 2026

49%
of US adults use AI chatbots, up from 33% in 2024
Pew Research, Jun 2026

AI-referred visitors to travel sites still convert 28% less often than other visitors, according to Adobe, but that gap has narrowed sharply since late 2024.

AI Summaries Reduce Clicks to Websites

Pew Research Center found that Google users who encountered an AI summary were less likely to click on links to other websites than users who did not see one, and very rarely clicked on the sources cited in the summary.

Google users are less likely to click on links when an AI summary appears
Share of visits to Google search pages with a click, March 2025

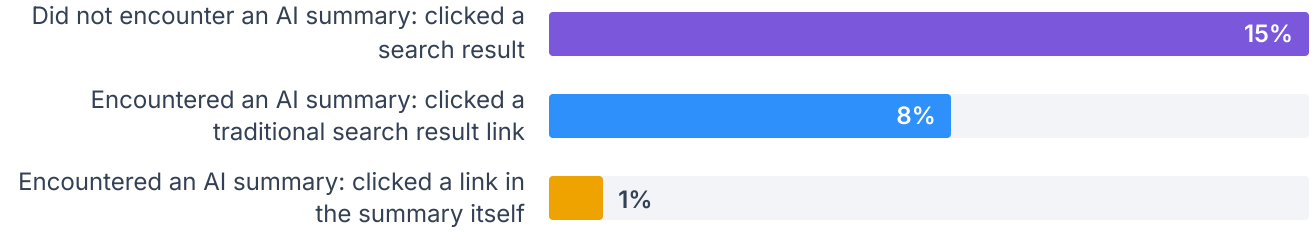

Pew Research Center July 22, 2025. Browsing data from 900 U.S. adults, March 2025.

Finding 01

Being Known by AI vs. Being Named or Cited by AI

AI Visibility measures the share of answers that name a property or cite its website. The most revealing distinction is between branded and unbranded prompts. Branded prompts supply the property's name, so visibility is near-certain. Unbranded prompts are where engines decide which properties to name or cite, and where most of the variation between properties occurs.

Branded prompt

"Does The Juniper House* in Asheville allow dogs?"

The traveler has already chosen the property. AI names it in 100% of answers.

Unbranded prompt

"What's the best pet-friendly hotel in Asheville?"

The traveler is still deciding. The median property is named or cited in only 55% of answers.

*The Juniper House in Asheville is a fictional name used to mask real property details. Results reflect real benchmark data.

AI Visibility on Branded vs Unbranded Prompts



Milestone Research — Medians of location-level rates, 1,600 hospitality locations. Hotels and resorts only: 54%.

Finding 02

Visibility Varies by Prompt Intent

Visibility depends heavily on the intent behind the prompt. Comparison prompts, which ask an engine to weigh competing hotels without naming the property itself, showed the highest visibility (75%), though on a small sample. Only unbranded prompts are counted in this chart.

Unbranded Visibility by Prompt Type

How often AI names or cites the property, by the kind of question a traveler asks without naming it

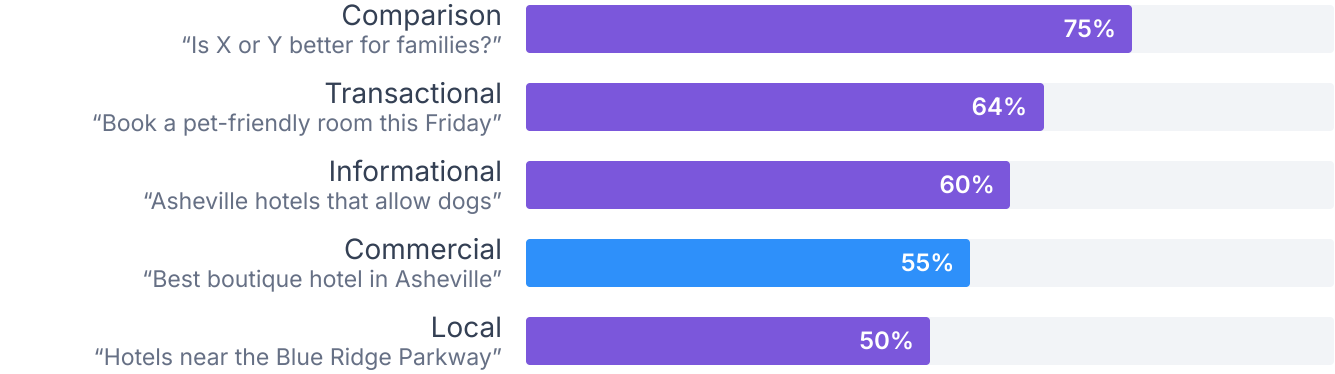

Milestone Research

Commercial prompts account for 73% of hospitality answers. Comparison is a small sample: 2% of hospitality answers, from 326 locations across 17 businesses (Low confidence). Treat 75% as a rough guide, not a market benchmark. Hotels only: 85%, from 39 locations across 15 businesses (also Low confidence).

What the Chart Is Saying

Comparison: highest, on a small sample (75%)

When a traveler compares named hotels, AI often brings in other hotels from the same competitive set. Properties in that set tend to appear.

Commercial: the largest volume

These prompts ("best boutique hotel…") make up 73% of hospitality answers, yet median AI Visibility (named or cited) is only 55%.

Local: lowest (50%; hotels only, 56%)

Depends on AI knowing where a property is and what is nearby. Properties with limited location information are named less often.

Local: the widest spread

From roughly 15% to 80%, with the top tenth near 95%. The gap suggests location content may be a differentiator and warrants further testing.

Implications for Hotels

Local prompts

Specific "near X" content: distances to landmarks and the airport, neighborhood pages and accurate map listings.

Commercial prompts

Reviews and OTA listings that reflect what the property is "best for": couples, pets, families, business.

Informational prompts

Plainly stated facts on the website: parking, EV charging, pet fees and accessibility.

Finding 03

Different AI Engines Trust Different Sources

ChatGPT, Gemini, Perplexity, Claude and Google AI Overviews run on different models and retrieval partners, and return different answers to the same prompt on the same day. For hotels and resorts, each engine has a distinct citation fingerprint: the mix of sources it relies on before naming a property.

Citation Footprint: Share of Cited Sources by Type

Hotels and resorts, by AI engine

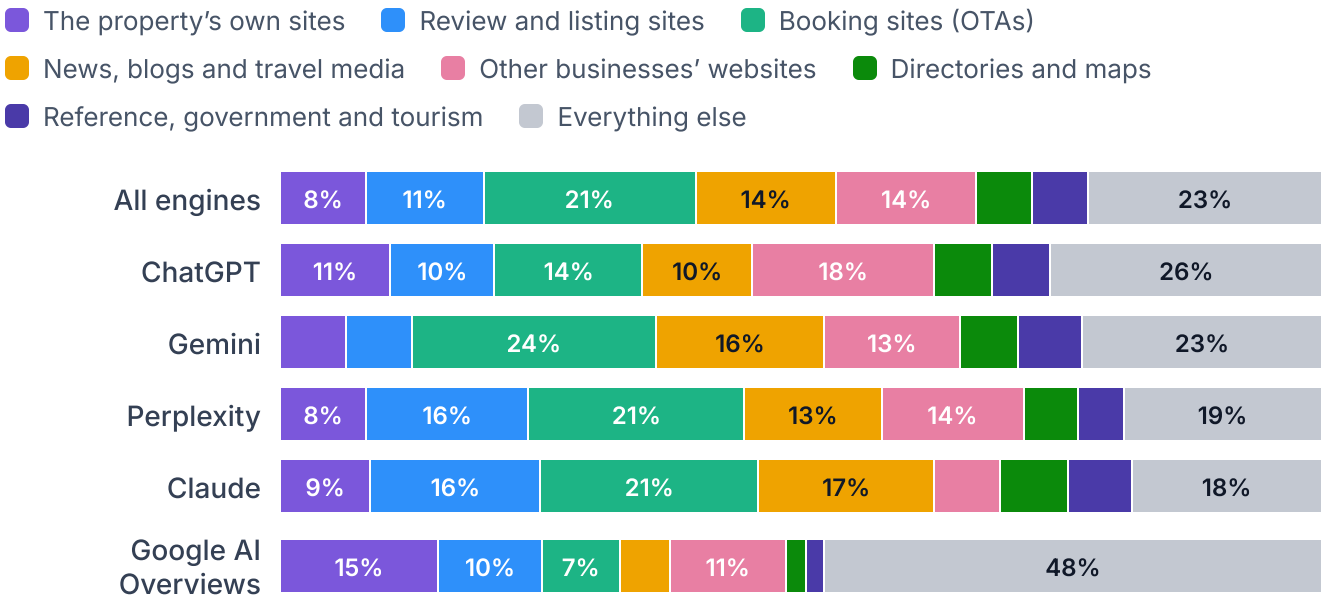

Milestone Research

Hotels and resorts, Q3 2026. Citations deduplicated per answer. Unlabeled segments are under 7%; "everything else" includes social media (about 2%).

Three patterns stand out. Booking sites are the most-cited source for hotels (21%), leading on Gemini, Perplexity and Claude. ChatGPT leans most on other local businesses' sites (18%). Google AI Overviews draws 15% of its citations from the hotel's own website, the highest of any engine, and cites booking sites least (7%).

How Each Engine Decides

Taken together, the fingerprints describe five distinct editorial models. Each draws on a different mix of sources, which may indicate where marketing effort is most visible.

ChatGPT
Other businesses 18%

Leans on other businesses' sites: local partners, venues and attractions. It cites hotels' own sites more than any other chat engine (11%) and booking sites least among them (14%).

Gemini
Booking sites 24%

The most booking-driven engine: 24% of citations from OTAs and 16% from travel media. It cites review sites least (6%).

Perplexity
Booking sites 21%

OTAs (21%) plus reviews (16%). Its hotel visibility rose from 65% to 70% this quarter, the largest significant rise of any engine.

Claude
Booking 21% · Media 17%

The broadest mix: OTAs (21%), travel media (17%, the highest of any engine) and reviews (16%, level with Perplexity).

Google AI Overviews
Own site 15%

Cites the hotel's own site most (15%) and OTAs least (7%), and names the fewest hotels per answer. Inclusion points the traveler to the property, not an intermediary.

Implications for Hotels

The five engines do not behave as a single channel. Gemini, Perplexity and Claude cite booking listings most. Claude and Perplexity also cite reviews more than the others, and Google AI Overviews cites owned pages most. Because engines differ, a single combined score can hide what drives visibility on each.

Finding 04

Booking and Review Sites Shape the AI Answer

This is one of the clearest patterns in the benchmark. AI engines build answers from the sources they retrieve and cite, and for hotels those sources are led by booking platforms, followed by review sites. Expedia, Tripadvisor and Booking.com are each cited in 25–29% of hotel answers. Owned websites are cited, but they form a small share of the evidence behind the answer.

49%

of unbranded prompt answers name or cite an OTA, for the typical hotel

Source: Milestone Research, hotels and resorts, Jul–Sep 2026

32%

of hotel citations from booking and review sites

Source: Milestone Research, hotels and resorts, Jul–Sep 2026

45%

owned-site citation rate: answers citing the property's own website, typical hotel

Source: Milestone Research, hotels and resorts, Jul–Sep 2026

44%

median share of voice for hotels: share of mentions against tracked competitors

Source: Milestone Research, hotels and resorts, Jul–Sep 2026

The Public Sites Engines Lean On

Share of hotel and resort answers citing each domain

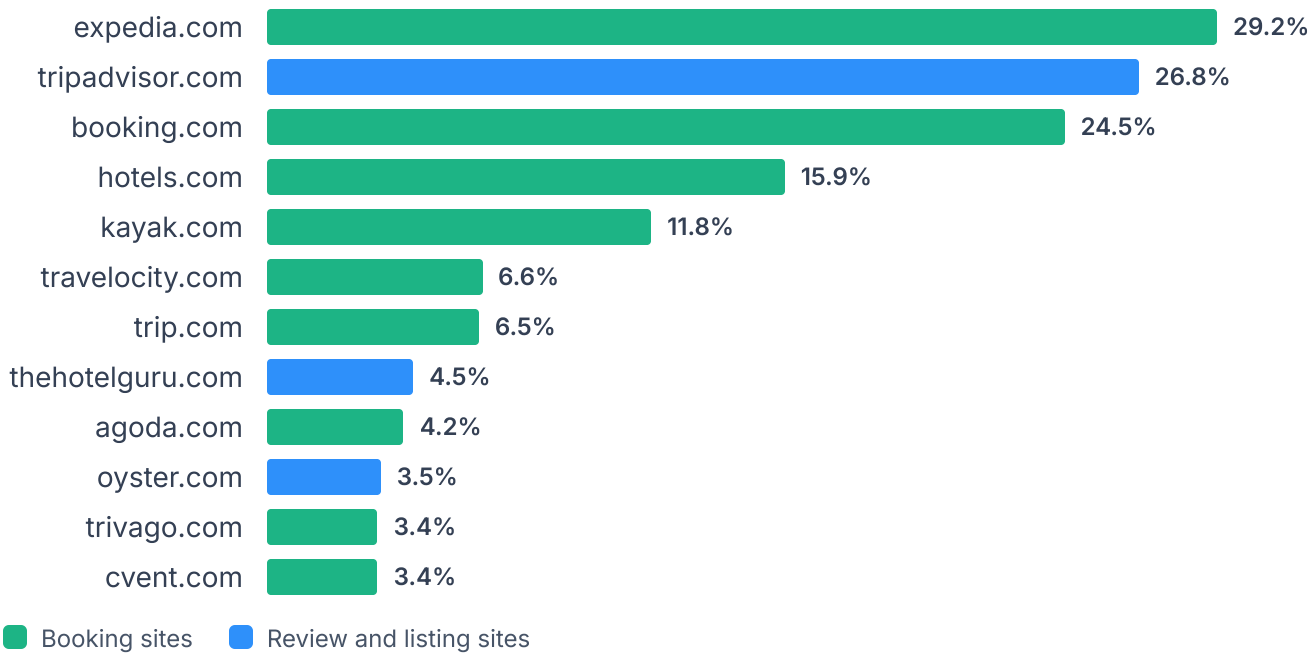

Milestone Research

Hotels and resorts only. Customer and hotel-brand websites are excluded from this ranking. By engine, Tripadvisor leads on Claude (52%), Perplexity (50%) and ChatGPT (27%); Expedia leads on Gemini (46%); Google AI Overviews cites any one domain rarely (Tripadvisor 7%).

Finding 05

Citing OTA Listings and Reviews Coincides with Higher Visibility

When an AI answer draws on a booking listing or a review page, Milestone observed higher AI Visibility for the hotel (named or cited): 65% vs 56% when no OTA listing is cited. The lift from review pages is smaller (+5 points). When the answer draws on a blog, guide or listicle, the hotel is named or cited less often, most of all with listicles (−10 points), as those articles tend to favor the hotels they already rank.

How Often the Property Is Named, by the Pages AI Cites

Hotels and resorts only. Share of answers that name or cite the property when the answer cites each page type, compared with answers that don't

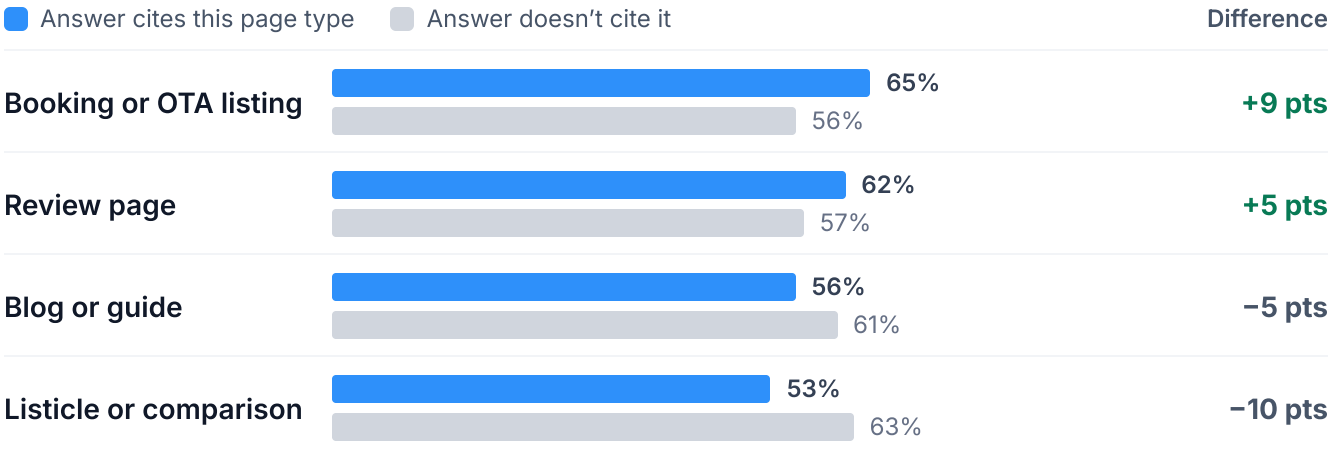

Milestone Research An association across the same locations, not proof of cause.

Most of What Shapes the Answer Sits Outside the Owned Website

For the median hotel, 45% of all answers link to the property's own site. Yet owned sites represent only 8% of hotel citations. The remainder comes from booking sites, review sites, travel media, and other local businesses.

Being Named Means Competing with Other Hotels in the Same Answer

When engines name a property, sentiment is overwhelmingly favorable: for the median hotel, 94% of mentions are positive and fewer than 1% negative. The greater challenge is competition within the answer. Unbranded prompt answers typically name several properties, and the median hotel holds a 44% share of voice against the competitors it tracks.

Finding 06

Perplexity Named Hotels More Often This Quarter (65% to 70%), Google AI Overviews Less Often (57% to 52%)

A benchmark establishes position; trend tests show how that position is changing. Milestone tested 138 potential changes between July and September on identical prompts and retained only those meeting every criterion: statistically significant after correction, at least 2 points in size, and consistent across most of the market.

Hotel and Resort AI Visibility, July vs September 2026

Identical prompts and locations, by engine

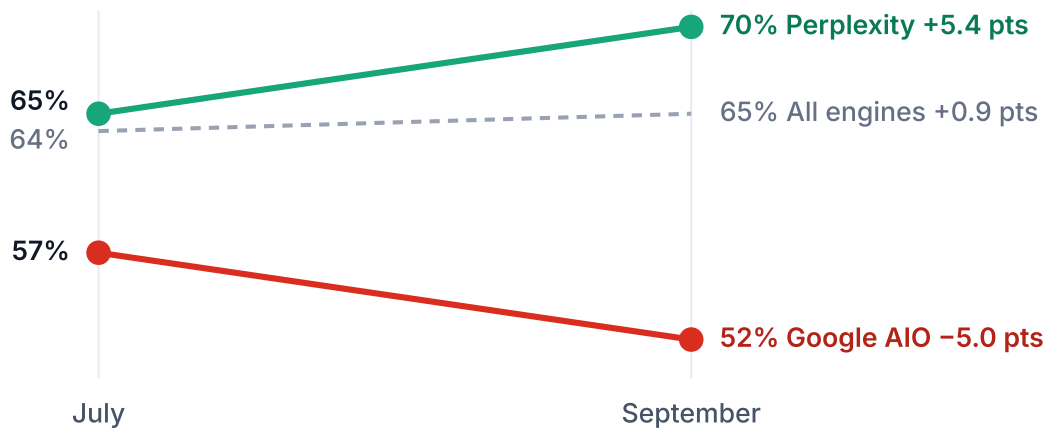

Milestone Research

Hotels and resorts only. Perplexity 65% to 70%; Google AIO 57% to 52%; all engines 64% to 65%, holding steady below Milestone's 2-point threshold.

72%

of hotel brands moved up on Perplexity

Source: Milestone Research, hotels and resorts, Jul–Sep 2026

97%

of hotel brands saw Google's AI hotel lists get shorter

Source: Milestone Research, hotels and resorts, Jul–Sep 2026

Google Is Shortening Its Lists

Google AI Overviews already returns the shortest answers of any engine, averaging 2.19 hotels per unbranded answer across the quarter, against 4.83 to 6.44 for the others. This quarter they contracted further: the number of hotels named in a typical unbranded prompt answer fell by more than a quarter (29%).

Hotels Named per Google AI Overviews Answer

Average brands named per unbranded hotel answer

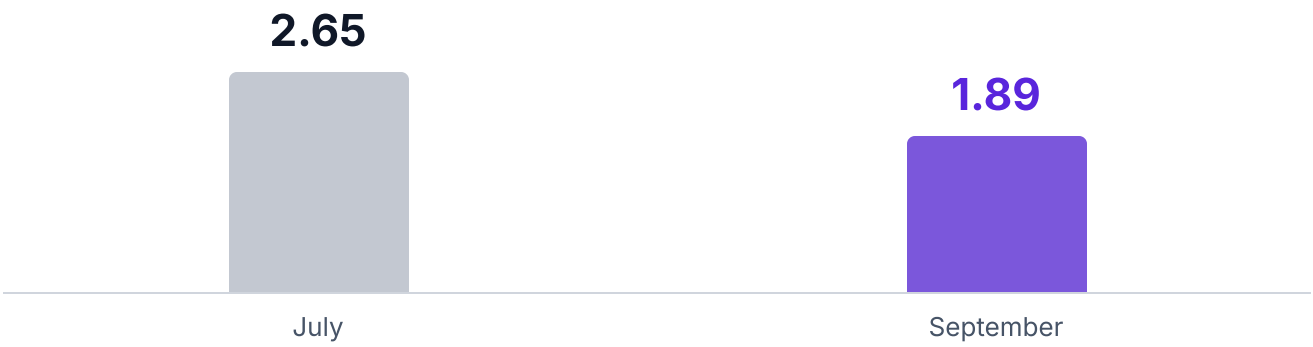

Milestone Research Identical prompts and locations, July vs September 2026.

Google AI Overviews now names fewer hotels

Google AI Overviews is listing fewer hotels per answer across the market. This is a market-wide pattern consistent with shorter answers; it does not rule out property-level content as a factor in any individual decline.

Perplexity names more hotels per answer

Perplexity now names 5.99 hotels per unbranded answer, second only to Claude (6.44). Monitoring limited to a single engine would not capture these differences.

Finding 07

No Clear Brand-Affiliation Advantage Was Established in This Sample

Brand-affiliated hotels were named or cited in 47% of unbranded prompt answers, independents in 55%. The ranges overlap, so the data did not establish a reliable difference between segments.

AI Visibility by Hotel Segment

Median AI Visibility, hotels only

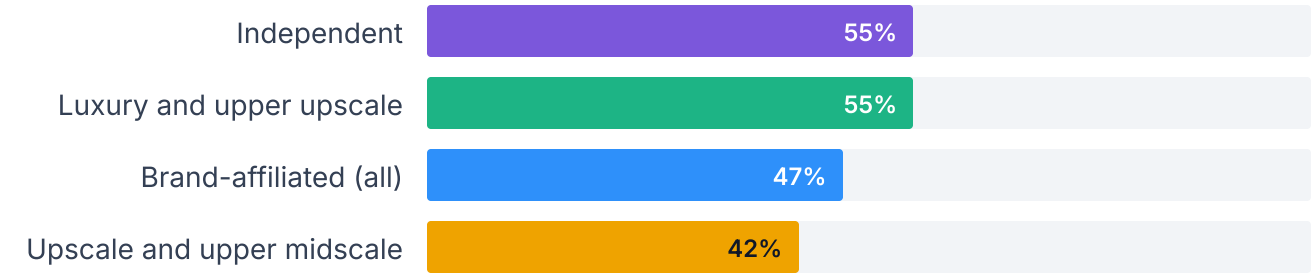

Milestone Research

Brand-affiliated: hotels whose name matches a known chain brand. Independent: no brand match. Chain-scale bands (luxury/upper upscale, upscale/upper midscale) are Milestone's approximate mapping of brand-affiliated hotels, so they overlap with "Brand-affiliated (all)". A franchised hotel whose name omits the brand is counted as independent. The midscale and economy band is not shown because the sample is too small.

What the Chart Is Saying

Independents (55%) and luxury and upper-upscale hotels (55%) had the same median; upscale and upper-midscale hotels were lowest (42%). Because the ranges overlap, these differences are directional only. One possible explanation, not tested here, is that some branded properties rely on brand-level pages with fewer property-specific facts.

Recommendations

Five Priorities for Hospitality Marketers

Each priority is grounded in the findings above and concerns measurement or content practice.

1 Map the Questions Travelers Ask

Map the questions travelers ask about properties like yours, including the follow-up queries AI engines generate. Track them as part of your visibility measurement.

2 Track Every Engine Separately

Each engine relies on different sources, and a blended score can obscure engine-specific differences. Include Perplexity in the tracking strategy: its hotel visibility rose from 65% to 70% this quarter.

3 Optimize Your Citation Sources

Booking and review sites account for 32% of hotel citations, with booking sites alone at 21%. Keep OTA listings complete and consistent, earn travel-media coverage and respond to reviews.

4 Strengthen the Owned Website as a Source

Publish factually accurate content on every key page to help engines represent the property accurately. Maintain a single source of truth and keep details consistent across channels. One proposed explanation is that clear, structured facts make it easier for AI engines to interpret a property's information (Milestone, "From strings to things"); this benchmark does not test that effect directly.

5 Test Content Formats by Intent and Engine

The benchmark associates booking listings, and to a lesser extent review pages, with higher visibility, and blogs, guides and listicles with lower visibility. Test formats against specific prompt intents and engines before investing, rather than relying on citation frequency alone.

Research Methodology

This edition of the quarterly Milestone benchmark, published in November 2026, analyzes AI answers recorded for Milestone hospitality customers between July 1 and September 30, 2026. All answers were captured during routine monitoring; none were re-run or simulated for this report. Findings reflect this customer base rather than the hospitality market as a whole and are reported in aggregate.

Sample

Group	Locations	Businesses	Share of answers
Hotels and resorts	1,128	115	46%
RV parks, campgrounds and communities	472	Fewer than 10	54%
Total hospitality	**1,600**	**119***	**100%**

*Some businesses own both hotels and RV parks, so business counts overlap. A *location* is one property, measured as its own profile. A *business* is the customer account that owns it. An *answer* is one engine's reply to one prompt on one day: 1,381,203 answers with text in Q3 2026. RV parks are 30% of locations but 54% of answers because they run more often (a median of about 1,470 answers per location, against about 320 for hotels). Citation, domain, sentiment, share-of-voice and trend findings use hotels and resorts only (1,128 locations). Location benchmarks cover all 1,600 hospitality locations, with hotel-only figures shown alongside. Comparison results are Low confidence (fewer than 20 businesses). Prompts are in English; 86% of locations are in the US.

Engines and collection

Engine	How it searches	Model in Q3
ChatGPT	Searches the web on every prompt	gpt-5-search-api
Gemini	Google Search grounding turned on	gemini-2.5-pro
Perplexity	Search built into the model	sonar-pro
Claude	Anthropic's web search tool	claude-sonnet-4-6
Google AI Overviews	Live Google search from the property's location	Google's own

Engines are queried through their APIs with live web search on. Each location runs daily, weekly or monthly, and each run asks every engine each prompt once. The median location was measured on 3 run dates with 38 prompts. 615 locations used all five engines and 968 used four, mostly without Claude.

Prompts and prompt types

Prompt libraries are drafted when a location is set up, from the property's website, name and area; customers can edit them. A prompt is *branded* when it contains the property's name or an alias; all others are *unbranded* (discovery) prompts, about 92% of answers. Every prompt carries equal weight, regardless of search volume.

Type	Example (illustrative)	Hospitality	Hotels	RV parks
Commercial	Best boutique hotel in Asheville	73%	77%	72%
Informational	Asheville hotels that allow dogs	13%	12%	13%
Transactional	Book a pet-friendly room this Friday	5%	4%	7%
Local	Hotels near the Blue Ridge Parkway	3%	2%	3%
Navigational	Juniper House phone number	3%	4%	2%
Comparison	Is Hotel A or Hotel B better for families?	2%	1%	3%

Share of answers by prompt type. Results use unbranded prompts only. Local: 50% (412 locations, 56 businesses), hotels 56% (209, 54), high confidence. Comparison: 75% (326, 17), hotels 85% (39, 15), low confidence; excludes comparisons naming the property.

Metrics

Measure	Definition
AI Visibility	Share of answers that name the property or cite a site it owns (5% of matches are citation-only). Measures inclusion only; whether an engine recommended the property was not assessed separately
Owned-site citation rate	Share of answers citing a website owned by the property's business
Booking-site presence	Share of unbranded answers naming or citing an OTA or booking site (hotels only)
Citation share by source type	Each of nine source types' share of cited sources, counted once per answer
Share of voice	Property mentions as a share of its own plus chosen competitors' mentions

What was excluded

Excluded or handled differently	Q3 count
Locations with fewer than 10 distinct prompts	74 locations
Internal, demo and test accounts; duplicate answers	14 locations; duplicates removed
Google AI Overviews results with no overview	65,220, left out of every measure
Empty answers from the other four engines	33,882, counted as not visible

Benchmark bands and publishing rules

One score per location for the quarter, read at the 25th, 50th (benchmark), 75th and 90th percentiles, with no business above 15% of a group's weight. Medians carry 95% ranges from 400 redraws; overlapping ranges are not reported as differences. Groups need 25+ locations and 5+ businesses.

Confidence	Businesses behind the number	Typical precision of the median
High	40 or more	About 6 points or better
Medium	20 to 39	About 10 points
Low	5 to 19	Wider than 10 points; a rough guide

Trend tests

Milestone pre-registered 138 tests (101 hospitality), comparing July with September on identical prompts, engines and locations, with 2,000 redraws and a Benjamini-Hochberg correction at 10%. A trend must move 2+ points, in the same direction for 60%+ of businesses, and survive removing the largest. Model settings did not change, so shifts most likely come from the engines.

Limitations

Figures are directional. Prompts are written for each property's own market, so rates likely run higher than a neutral sample. Results reflect API answers, not signed-in app sessions, and associations do not show cause. Accuracy, clicks and bookings are not measured; no customer is identifiable.

Sources

Primary source: Milestone AI Visibility dataset

1,381,203 live AI answers collected for 1,600 Milestone hospitality customer locations, July to September 2026. All Milestone figures in this report come from this dataset.

Related reading

"Why entity authority is the foundation of AI search visibility" Search Engine Land, Mar 16, 2026

"Top tips to create AI visibility for your hotel" Milestone webinar, Aug 21, 2026

"What it takes for AI to find, trust, and act on your brand" AI Hospitality Alliance, Sep 5, 2026

Secondary research: Phocuswright (March and July 2026), Adobe Digital Insights (June 2026), Pew Research Center (July 2025 and June 2026).

About the Partners

Milestone partners with leading voices in hospitality AI, distribution and technology - the AI Hospitality Alliance, HEDNA and Aven Hospitality - to bring Milestone's research to the wider industry, giving it a common, evidence-based view of how AI engines discover, name and cite properties.

Milestone Inc. Research lead and data source

Milestone is the AI-native digital experience platform that turns AI visibility into incremental revenue and leads. Unlike platforms that simply measure visibility, Milestone combines AI search intelligence, content optimization, AI-first websites, and conversion technology to help brands get discovered, engage customers, and drive measurable business growth. Brands that move to Milestone have seen 10–25% more direct revenue (company-reported; not measured in this study) by improving visibility across AI and traditional search, closing gaps in content and schema, and delivering high-converting digital experiences on agent-ready infrastructure.

www.milestoneinternet.com

AI Hospitality Alliance (AIHA) Release partner

The AI Hospitality Alliance is an independent, vendor-neutral organization bringing together hotel leaders, technology providers, startups, investors and academics to advance responsible AI adoption across hospitality. Through collaborative research, practical education, industry frameworks and events, AIHA helps hospitality professionals understand emerging technologies, evaluate opportunities and turn innovation into action. Its work gives the industry a collective voice in shaping how AI transforms hotel operations, distribution and the guest experience.

www.aihospitalityalliance.com

HEDNA Release partner

HEDNA is a global community dedicated to making hotel distribution more effective through collaboration, education and knowledge sharing. Bringing together hotel companies, technology providers, intermediaries and industry experts, HEDNA helps members navigate evolving technologies and commercial practices. Its conferences, research and practitioner-led working groups address challenges across connectivity, content distribution, revenue, data and payments, providing practical guidance and professional connections that help move the hospitality industry forward.

www.hedna.org

Aven Hospitality Co-sponsor

Aven Hospitality is an independent global technology and SaaS leader in the hospitality industry, powering how hotels sell, distribute, and deliver guest experiences. SynXis, the leading global hospitality commerce and distribution platform, coordinates hotel information across 600+ integrations to ensure accuracy, consistency, and competitiveness wherever travelers discover, compare, and book. Trusted by 35,000+ hotels across 190+ countries, Aven Hospitality serves as the connective layer behind many of the world's most iconic brands, simplifying complexity behind the scenes so hoteliers can operate with confidence today and stay ready for what's next.

www.avenhospitality.com