Milestone

Q3 2026

From Ranking to Being Cited:

The Execution Plan for AI Search Visibility in Hospitality

By Anil Aggarwal, CEO, Milestone, Inc.

Published September 2026

milestoneinternet.com

From Ranking to Being Cited: The Execution Plan for AI Search Visibility in Hospitality

By Anil Aggarwal, CEO, Milestone Inc.

Milestone CEO Anil Aggarwal outlines a four-step execution loop for hotels to build visibility in AI search, drawing on data from 3,500 hospitality websites.

- AI referral traffic across the 3,500-site panel has grown exponentially in sixteen months, and AI usage for travel planning now exceeds 76% among Gen Z and Millennial travelers (Skift Research).
- The guest journey has forked into three paths — click-through to the website, discovery and booking inside the LLM, and agent-led transactions — each rewarding different investments.
- The execution loop runs four steps continuously: detect visibility gaps at the prompt level, generate content that closes them, deploy schema, and publish across listings and indexing protocols.

Why it matters for hospitality
AI engines recommend hotels based on structured, verifiable, current data rather than traditional rankings, and every grounding failure hands the citation and the booking to an OTA charging 15–25% commission. The four-step loop gives hotel teams a repeatable operating discipline for AI visibility, and the 90-day plan makes it actionable regardless of property size or existing tech stack.

Every hotel executive we talk to has now accepted the premise: guests are using AI to plan travel. The question has moved past "is this real?" to something harder, "what exactly do we do on Monday morning?" That is a fair question, and most of the industry conversation has not answered it well. There is plenty of commentary about the arrival of AI search and very little about the operating plan. What follows is that plan, drawn from what we see across 3,500 hospitality and multi-location websites running on our platform.

1. Start with what the data says

Two numbers frame the urgency.

First, adoption. Skift Research's 2025 U.S. travel trends work puts AI usage for travel planning above 76% among Gen Z and Millennial travelers. These are not fringe users. They are the guests who will fill your rooms for the next twenty years.

Second, traffic. Across our 3,500-site panel, referral sessions from AI assistants have grown exponentially in sixteen months. ChatGPT accounts for the overwhelming share and grew roughly fifteen-fold in that window; the true totals are likely higher still, since Google does not break out Gemini referrals. The growth curve is reminiscent of what we saw with mobile traffic in 2010-12.

Bottom line: your guests are already using AI engines to plan travel, find hotels, and book and that number will only grow.

2. The customer journey now has three shapes, not one

The reason a single "AI strategy" is hard to write is that AI has not replaced the funnel. It has forked it into three.

a. **Journey one: the guest still arrives.** They start in Google or an LLM, see a citation or a result, click, and land on your website. This is the familiar path, and it still carries the most volume.
b. **Journey two: the guest never leaves.** They stay inside ChatGPT or Gemini through discovery, comparison, and increasingly through booking itself. Your brand either appears in the answer or it does not exist for that guest.
c. **Journey three: an agent transacts on the guest's behalf.** A personal AI agent queries your rates and availability directly and completes the booking. The agent never loads a page. It never sees your hero image.



Figure 1: How Agentic Commerce is changing the guest journey

Your website is the foundational data hub for all three. But only journey one rewards the things hotels have historically optimized - design, imagery, page speed as a user experience. Journeys two and three reward something else entirely: structured, retrievable, verifiable data.

3. Why the mechanics dictate the plan

To build the right plan you have to understand what happens between a guest's question and an AI's answer. The sequence is roughly: **understand → fan out → retrieve → filter → ground → cite.** Two of those steps matter most operationally.

Query fan-out. Guests stopped typing keywords. They now describe situations. "I'm going to Chicago for a conference. I need a hotel with a good desk to work at. Would love a restaurant, spa, and gym. Prefer walking distance to the Magnificent Mile." The model silently decomposes that into a dozen sub-queries - location, workspace, dining, wellness, walkability, price - and hunts for answers to each. If your site answers all twelve, you surface. If it answers three, you do not.

Grounding. Models are increasingly unwilling to assert facts they cannot source. Before recommending a property, the model looks for something solid to lean on: your website, your listings, your structured data. If your information is clear, complete, and current, it cites you. If it is thin or stale, it leans on someone else - and that someone else is almost always an OTA. Every grounding failure on your site is a citation handed to a distribution partner who charges you 15 to 25 points for it.

4. The execution loop

The plan that works is a loop, not a project. Four steps, run continuously.

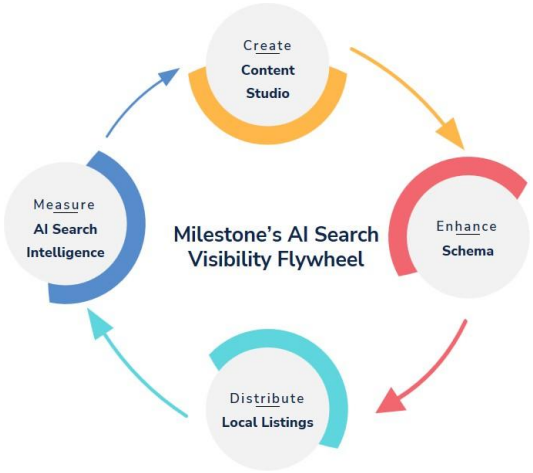


Figure 2: Milestone's AI Search Visibility Flywheel Platform

1. Measure – AI Search Intelligence

You cannot fix what you cannot see, and traditional rank tracking is useless here - there are no positions to track. What you need is prompt-level visibility measurement: for the 40 to 50 prompts that actually matter to your property, how often does your brand appear in ChatGPT, Gemini, Perplexity, and Claude? With what sentiment? And which sources are being cited instead of you?

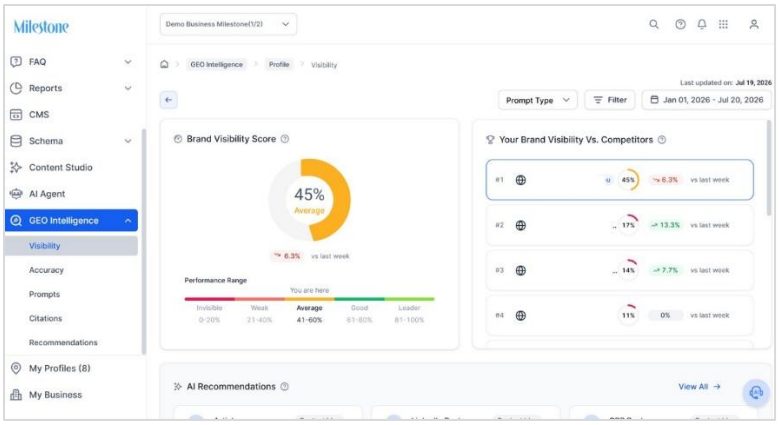


Figure 3: AI visibility for your hotel vs. competitors, across engines

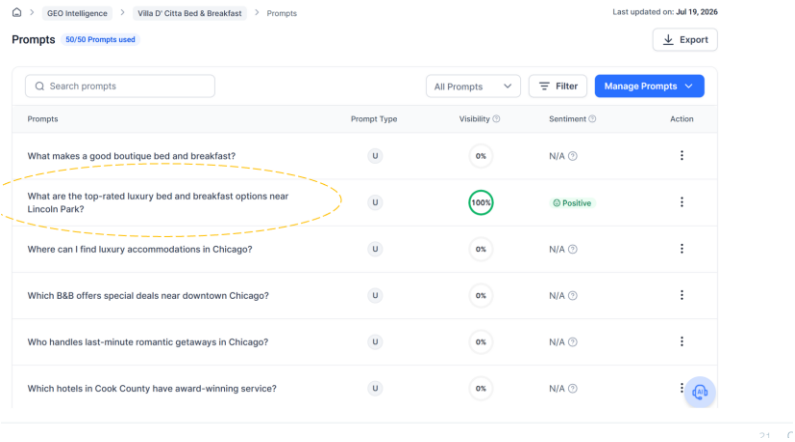


Figure 4: Visibility measured prompt by prompt for your property

Milestone's AI Search Intelligence platform offers the visibility for hotels and brands on AI Search LLMs. The typical first read is sobering. A property we measured recently sat at 37.5% overall AI visibility and 19% on non-branded prompts - meaning it showed up when someone asked for it by name and largely vanished when they described a need. Eleven of 48 prompts scored 100%. Nineteen scored zero. The zeros are the roadmap.

Also audit brand accuracy across your listings and profiles. AI engines cross-check. Conflicting hours, addresses, or amenity data across Google, Apple, Bing, and the aggregators actively suppresses you, because inconsistency reads as unreliability.

2. Create: Generate the content that closes the gap – Content Studio

Each zero-scoring prompt is a content brief. "No page targets 'best hotel near Napa wineries for couples'" is not an abstract insight - it is next month's article.

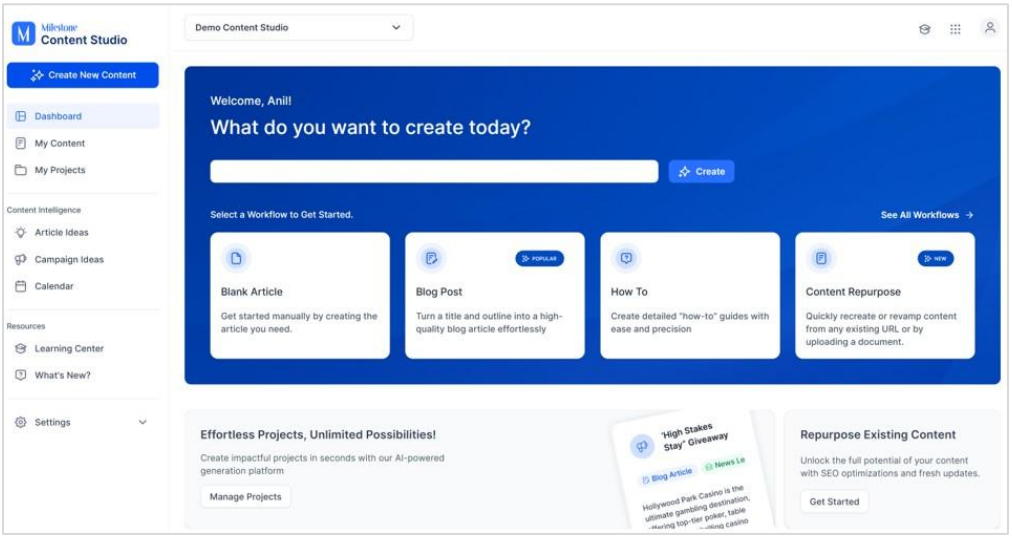


Figure 5: Dashboard view of Milestone's Content Studio platform for automating content creation workflow

The volume required is the hard part. Closing 19 gaps at hotel marketing department speed takes a year. This is where generative content platforms earn their place, provided three conditions hold: output is tuned to your brand voice and target personas, claims are specific and traceable to named sources, and human reviews before publication. Generic AI-written content fails at the grounding step for the same reason thin content does. Specificity is the whole game - named sources, real data, verifiable claims.

Freshness matters independently. FAQs, events, offers, and image libraries that update on a cadence, signal for an active, maintained source. Our events data is a useful proxy: properties running structured local event content saw an average 36% monthly lift in impressions and 41% in pageviews, with measurable revenue attribution to those pages.

3. Enhance: Deploy Schema

If there is one technical investment to prioritize, it is structured data. Schema is the translation layer between how you describe your hotel and how a machine files it.

"We're the perfect hotel for executive meetings near DFW, with private boardrooms, airport convenience, and flexible meeting space" is human language. The machine needs: Business Type: Hotel. Audience: Business Travelers, Meeting Planners. Amenities: Boardroom, Meeting Space,

Restaurant. Location: Near Airport. Use Case: Executive Meetings, Leadership Retreats, Corporate Events.

That is what gets stored, connected, and retrieved later. Schema across your organization, individual properties, offers, events, and reviews is how you get represented as an entity in a knowledge graph rather than as a page of prose that may or may not be parsed correctly.

4. Publish and distribute – CMS and Local Listings

Content that sits in a staging queue earns nothing. The loop only compounds if publication is fast and distribution is broad - production site, local listings across Google, Apple, and Bing, and real-time indexing protocols like IndexNow so changes are discovered in minutes rather than weeks.

5. Then get ready for the agents

Journeys one and two are addressable with content and schema. Journey three is not.

When an AI agent books on a guest's behalf, it queries your data directly through emerging protocols - MCP, Universal Commerce Protocol, AP2, NLWeb. It needs clean rates, live availability, and structured content exposed through a callable endpoint. Brands that have that plumbing are bookable by agents. Brands that do not are, functionally, invisible to them.

Most hotels are not ready, and that is fine for the moment. But this belongs in the 2027 budget conversation, not the 2029 one. Ask your website and booking engine vendors what their protocol roadmap is. The answers will be clarifying.

6. A realistic first 90 days

Days 1–30: Measure. Establish your AI visibility baseline across engines, define the 40 to 50 prompts that matter, audit brand accuracy across listings, and identify your top competing citation sources.

Days 31–60: Fix the foundation. Deploy or repair schema across organization, location, offers, events, and reviews. Resolve listing inconsistencies. Confirm AI crawlers are not blocked - a surprising number of hotel sites are quietly excluding them at the robots.txt or WAF layer.

Days 61–90: Build the content engine. Turn your zero-visibility prompts into a twelve-month editorial calendar, establish the brand-voice and review workflow, and start publishing against the gaps.

Then re-measure and run the loop again.

For twenty years, hospitality digital strategy was about being findable. The new bar is being *citable* - being the source an AI trusts enough to name when a guest describes exactly what they want. That

is not a marketing campaign. It is an operating discipline, and the properties that build it now will hold that position for a long time.

Author, Anil Aggarwal, is the CEO of Milestone and is a renowned speaker and thought leader in AI search visibility. He has authored numerous articles and speaks regularly at major hospitality conferences.

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 generate 10–25% more direct revenue by improving visibility across AI and traditional search, closing gaps in content and schema, and delivering high-converting digital experiences on agent-ready infrastructure.

Founded in 1997 and headquartered in Silicon Valley, Milestone's DXP powers 3,500+ websites for hotels and multi-location brands, including IHG, Marriott, and Sun Outdoors. Its AI Visibility Platform is used by hundreds of additional websites to improve visibility in AI and traditional search.