

DISCOVERY · AI VISIBILITY

Be the hotel the AI recommends. Booked direct.

Discovery gives your hotel a machine-readable, citable page — one that ChatGPT, Claude, Perplexity and Gemini read and quote on their own, with your booking link. And it comes with proof: every week it measures all four engines — whether they mention you, how they describe you, where they send the guest, and exactly what to fix.

A readable page · Four engines, weekly · A closed loop

PUBLISHED BY
Assist Intelligence Kft.
peaqplus.com ·
sales@peaqplus.com

TRAVELLERS NOW ASK AN AI

The new front door is an AI assistant.

When someone asks an assistant "a wellness hotel near the lake I can book directly?", it doesn't show ads. It reads the web, picks a few trustworthy sources, and answers. If your hotel isn't readable in a way it trusts, you're simply not in the answer.

Discovery makes your hotel that trustworthy source. A clean, structured page the assistants can quote — your name, your live rate, your score, your location — so when the question comes up, your hotel is the one they cite, with a link that goes straight to your booking engine.

Not an ad you pay to place —
[a citation you earn.](#)

WHY IT WORKS

- **Answer-engine ready.** ChatGPT, Claude, Perplexity, Gemini and Copilot all read the open web — your Discovery page is a readable, quotable source for every one of them.
- **Cited, not advertised.** A source the assistants quote on their own. Not a paid slot — you can't buy your way into the answer.
- **The link goes to you.** The link in the answer is your booking engine, not an OTA. You earn the placement by being readable — not by outbidding.

ONE READ, THE WHOLE HOTEL

A profile the machines understand.

An assistant has a fraction of a second to decide whether your hotel fits the question. Discovery gives it everything in a single read — and the numbers never go stale: the rate and score refresh every day, so the agent reads a fresh storefront, not last year's brochure.

Live direct rate

REFRESHED DAILY

The real, bookable price — straight from your own system, not a marketplace list rate.

Guest score

WITH ITS SOURCE

The score with its source named, so the agent trusts it enough to cite it.

The full entity

YOUR WORDS

Address, amenities, your own guest FAQ, hero image, official links — edited in-app, in your language.

Direct-booking URL

BUILT IN

The next step, built into the page — every path ends at your own booking engine.

Every link goes to your engine, never a marketplace —
the guest, and the margin, are yours.

When a traveller arrives from an AI answer and books, it's a direct reservation: no OTA commission, and the click carries through to your own analytics — so you see the demand the assistants send you, not guess at it.

MEASURED WEEKLY, ACROSS FOUR ENGINES

Not just "are you mentioned." The whole picture.

Every week the same frozen panel of real traveller questions goes to Perplexity, ChatGPT, Claude and Gemini — in English and Hungarian, with your hotel's name never in the question. Then it turns the answers into four things you can act on.

01 — FOUR ENGINES + SHARE OF VOICE

Who mentions you — and how loud

All four engines, search-based (Perplexity) and training-knowledge (ChatGPT, Claude, Gemini) measured apart — plus your Share of Voice against the compset you actually compete with.

02 — SOURCE DOMAINS

Where the AI sends your guest

Every cited source classified — your own site, a commission OTA (Booking, Szállás.hu, Trivago...), or other. The direct-vs-OTA split, made concrete.

03 — AI MIRROR

How the AI describes you

Not how often — how. The sentiment, strengths and positioning it paints of you, with verbatim quotes in the AI's own words. Catch a stale claim before a guest reads it.

04 — TWO AUDITS + CLOSED LOOP

What to fix, then fixed

A website AI-readability score (0–100 + grade) and a content score on experience, expertise, authority, trust — turned into FAQ suggestions and re-measured next week.

How it stays honest: a frozen panel (a moving number is a real trend, not the noise of changing questions), your name never in a question, and answers scored by deterministic code — no AI judging an AI. We ask the models directly; consumer apps that add live web search can answer differently, so the search-based run is the closest stand-in for a real traveller.

THE FINE PRINT

Visibility we measure. Not visibility we promise.

What Discovery *doesn't* do matters as much as what it does. We don't claim anything we don't measure.

- **Not a paid slot.** You earn the placement by being readable — you can't buy your way into the answer.
- **It doesn't promise a number.** Your baseline is a starting line, not a guarantee — content work and time move it, and every weekly run shows whether it's moving.
- **Four engines, measured — not more than we check.** Perplexity, ChatGPT, Claude and Gemini. The search-based one moves first; the training-knowledge ones catch up on a slower clock.
- **It runs itself.** Switch it on and fill your content — the first measurement fires on its own, your questions regenerate when you edit, then a weekly run carries the trend. No button to press.

Pricing

Add-on to any plan **€49 / mo (intro)** Per property, flat — any hotel size **no setup fee**

On top of BI Core. The public page, the four-engine weekly measurement, both audits, AI Mirror and the closed-loop suggestions ship together in one price — AI-visibility tracking alone is sold elsewhere as its own subscription. Live today on stay.peaqplus.com with our first hotels.

See your own baseline.

A 45–60 minute walkthrough: a live Discovery page and a real weekly run — four engines, Share of Voice, the AI's own words about the hotel, both audits and the fix-list — plus the crawler log that proves the bots are reading.

sales@peaqplus.com