

IS HOTEL SAVOY IN CHATGPT?

Hotel Savoy Bern

MEASURED

30 September
2026

QUESTIONS

10 × 3 languages

ENGINES

4 grounded APIs

ANSWERS KEPT

116 of 120

SOMEWHERE IN SYDNEY, LAST TUESDAY

A couple planning four nights in Switzerland decide they want two of them in Bern, and rather than open Google and work through ten blue links they type a paragraph into an assistant and ask it where they should stay. Three hotels come back, and there is no second page.

That couple is not hypothetical, because Australians are 17% of your reviewers and the second largest group after Americans. So we asked the same systems the same kind of question thirty times across three languages, kept every answer, and this is what came back.

01 • WHAT THE MACHINES SAY TODAY

46.6%

You were named in 54 of 116 answers. A guest asking one of these thirty questions has a little under a coin flip's chance of hearing your name.



■ Named 54 ■ Not named 62

English 52% · German 37% · French 50%

What matters is that the rate is nowhere near evenly spread, because two of the questions name you every single time while another names you in no language at all, and the weakest of the three languages turns out to be the one your own website opens in.

02 • QUESTION BY QUESTION

You own two questions outright, and lose one completely.

Ten questions a guest actually types, each asked in English, German and French. The count is how many kept answers named you. A call the engines returned without any sources is discarded rather than counted.

THE QUESTION A GUEST ASKS			NAMED
01	where should I stay in Bern		5/12
02	best hotel in Bern old town		4/12
03	hotel near the station, walking distance		11/11
04	best hotel in Bern for a business trip		2/12
05	good four star hotel in Bern		12/12
06	nice hotel for a weekend with my partner		3/12
07	family friendly hotel with children		4/12
08	arriving by train, no car		8/10
09	hotel for a conference or event		0/11
10	quiet hotel in central Bern		5/12

TWO QUESTIONS YOU OWN COMPLETELY

“Good four star hotel in Bern” names you in 12 of 12 answers, and “hotel near the station, walking distance” in 11 of 11. Every engine, every language. Your location and your class are established facts as far as these systems are concerned, and no work is needed there.

ONE QUESTION YOU LOSE ENTIRELY

Conference and event: 0 of 11. Not one engine in any language named you, in the city that hosts the federal parliament. Business travel is nearly as bad at 2 of 12, and in German it is 0 of 4.

03 • DIAGNOSIS

They are not rejecting you. They are not considering you.

546

searches the four engines ran on their own behalf, from 30 guest questions

11

of 116 answers where any engine searched for “Hotel Savoy” by name

37%

German, against 52% English and 50% French

When a guest asks one question these systems quietly write and run a great many of their own, and we kept all 546 of them. They work in two steps, first casting around for candidates and then running a verification search on each hotel they are seriously considering, with the qualities they already associate with that hotel quoted straight into the query.

Across 116 answers, your name turned up in a verification search eleven times. For most of these questions you are not losing a comparison so much as never entering the shortlist where the comparison happens, which is a very different problem from a weak review score and it has a very different fix.

That pattern lines up neatly with the questions above, because wherever an engine already holds a firm fact about you such as your location or your star class you get named every time, while the questions that hinge on something nobody has published in a form a machine can read, whether that is meeting space or quietness or a room that sleeps a family, leave you out of the running altogether.

YOUR GUESTS ALREADY KNOW WHAT THE ENGINES DO NOT

“With the outer shades down, that room stayed pitch dark. Possibly the quietest city hotel I have ever been in.”

A guest from the United States, March 2025

“We booked 2 rooms for a family of 4, they gave us the connecting rooms which is very convenient.”

A guest from Canada, August 2024


“Catered for gluten free breakfast. Allowed early entry to room.”

A guest from Australia, June 2025

Now compare that to the machines. The German question for a quiet central hotel names you 0 of 4 times. The family question names you 4 of 12, and 0 of 4 in French. Quiet is real, connecting rooms are real, the dietary flexibility is real. Every one of those facts lives in a guest's prose and in none of the places a machine reads.

04 · WHO ELSE GETS NAMED

One comparison we can stand behind, and one we will not make yet.



What these answers do support. Hotel Schweizerhof turns up in 82 of the same 116 answers against your 54, and that particular comparison holds up because both names are distinctive enough to count without ambiguity, which means a

hotel in your own city is being named half again as often as you on exactly the same thirty questions.

What we are not going to claim. We will not hand you a ranking of Bern's hotels, because we asked thirty questions of our own choosing rather than taking an inventory of the city, so nothing here can tell you your position among all the hotels in Bern or among all the four star ones. Counting every name honestly would mean resolving “Savoy” and “Hotel Savoy Bern” back to one hotel across the sixty or so distinct names these answers contain, and until we have done that any league table we showed you would really just be a table of the hotels we happened to think of looking for.

We mention it because an earlier draft of this page did rank you, built on a list of competitor names we had written out from memory, and it quietly missed hotels that appear in more than a third of the answers. That felt worth telling you rather than tidying away, and it is the first thing we would fix, which is a day's work rather than a project.

Why the rest of the report is unaffected. Every other figure here is about your name and no one else's: whether an answer contains it, in which language, from which engine, and whether any engine searched for you by name. None of those depend on identifying your competitors correctly.

05 • REVENUE AT RISK

What the gap is worth.

Two of the three numbers below are measured. The third is yours, and we have not guessed it.

MEASURED

You are absent from 62 of 116 answers, which is 53.4%, and 2.2 nights is the average stay in your own published reviews.

YOURS TO SET


How many bookings a month reach you through an AI assisted search. Nobody publishes a credible figure for this and we will not invent one, so the table runs a range and you tell us which row is right.

BOOKINGS A MONTH YOU ARE NOT WINNING	AT CHF 200	AT CHF 240	AT CHF 280
10	52 800	63 400	73 900
25	132 000	158 400	184 800
50	264 000	316 800	369 600
CHF a year, calculated as bookings × 12 months × 2.2 nights × the rate. The three rates are placeholders until you give us your own average daily rate by segment, at which point this becomes one column instead of three.			

The conference question is where this concentrates. A meeting booking is a multi room, midweek, repeating booking, and you are named in none of the eleven answers we captured.

06 · WHERE THE FACTS COME FROM

The sources these engines actually read about Bern hotels.



Your own website is cited 21 times, which is essentially level with Bellevue Palace's 22. Your site is being read. The problem is what it does not say: it publishes no structured data at all, so everything on it has to be inferred from prose rather than read as fact.

A correction to our own opening hypothesis. Before measuring, we found three third party sites publishing machine readable hotel data about you with a wrong postcode, a phone number you do not answer, and a rate in euros. We expected those to be the cause. In this run the engines cited them zero times. They remain a risk worth fixing, and they are not what is costing you the conference question. We are telling you this because a report that quietly drops its own wrong guess is not worth paying for.

07 · WHAT WE WOULD DO ABOUT IT

Five moves, in the order the evidence puts them.

Every one of these comes out of a number in this report rather than out of a menu. None of them is about writing more marketing copy, and none of them touches what a guest writes about you.

1

State the facts you already have, in a form a machine can read

Your website publishes no structured data whatsoever, so nothing states the 64 rooms or the pedestrianised street or Bistro Savoy or the two minutes to the platform or the connecting rooms your guests keep thanking you for, which means everything an engine currently knows about you it has had to infer from your prose or borrow from a booking site.

Evidence: zero structured records on hotelsavoy-bern.ch, German and English. This is the floor, and it takes an afternoon.

2

Win the conference question

You appear in none of the eleven conference answers, in the city that hosts the federal parliament, and since a meeting booking tends to be multi room, midweek and repeating, this is the most valuable single gap in the report. The work itself is simply stating what you can actually host, in the places the systems answering that question are already looking.

Evidence: conference and event, 0 of 11 across three languages.

3

Close the German gap

German is your weakest language at 37% against English at 52%, which is an odd result for a hotel in a German speaking canton whose own site opens in German, and two of the three answers we threw away for citing nothing at all were the same engine working on German questions, which suggests the German language sources it was reaching for simply were not there to find.

Evidence: de 37%, en 52%, fr 50%. Business travel in German: 0 of 4.

4

Put your qualities where the engines actually read

TripAdvisor, Booking, hotels.com, bern.com and Expedia account for 425 of the citations behind these answers. Your own site accounts for 21. That ratio is not a problem to fix, it is a map: the words “quiet”, “connecting rooms” and “meeting space” need to appear in the places already being read, not only on a page of your own. This is earned coverage and profile work, and it is the slow expensive part.

Evidence: 3 414 citations recorded. Your site 21, Bellevue Palace's site 22.

5

Measure it weekly, the same questions on the same day

This report is one Tuesday, so on its own it tells you where you stand without telling you whether anything you change is working, whereas the same thirty questions asked every week become a line you can put in front of ownership and the only reliable way to tell a real improvement apart from the engines simply having a different day.

Evidence: every figure here is reproducible from a stored answer dated 30 September 2026.

What none of this is. It is not a rewrite of your website, because your website is the source these systems trust least. Every hotel's own site says that hotel is the best. They believe Booking, TripAdvisor, the local press and what your guests wrote, and getting those to agree about you is what takes months.

08 • WHAT THIS REPORT CANNOT TELL YOU

The honest limits, and why a conversation closes them.

Four of these we could close in an hour with you in the room, and one of them we cannot close at all, so it is written here rather than buried at the back.

We chose the questions, not you

These ten are our reading of how a guest looks for a hotel in Bern. You know which guests you actually want and which you make money on. An hour on your ideal guest, by segment and by season, would change several of these ten, and the changed ones are the ones worth measuring weekly.

Your rate is missing

The revenue table runs three placeholder rates because we do not have yours. Your average daily rate by segment turns that table from three columns into one number, and it is the number that decides whether any of this is worth your time.

Two of the six engines are not in these figures

Google AI Overviews has no interface we can query, and DeepSeek only searches inside its own app. Both have to be run by hand, and they are not in any percentage on this page. We would rather show you four measured properly than six measured loosely.

This is one day

30 September 2026, one pass. These systems are not stable week to week, so a single reading establishes a position and not a trend. The second reading is where the useful part starts.

And the biggest limit: we asked short questions, and your guests are typing paragraphs

We kept the ten questions short because short questions are comparable across languages, but real guests increasingly write several sentences at a time, and every extra sentence adds a condition you either satisfy somewhere on the record or quietly lose on. The engines are already doing this to themselves, since from our thirty questions they went on to write and run 546 of their own.

WHAT WE ASKED • QUIET HOTEL IN CENTRAL BERN

"We're flying into Zurich in March and want two nights in Bern. My wife is a light sleeper so somewhere genuinely quiet, walking distance from the station so we don't need a taxi, old town if possible, and ideally under CHF 300 with breakfast."

CONDITIONS IT CARRIES, AND WHETHER YOU STATE THEM ANYWHERE READABLE

station, 2 min

old town

genuinely quiet

rate with breakfast

light sleeper, blackout

WHAT WE ASKED • FAMILY FRIENDLY HOTEL IN BERN WITH CHILDREN

“Travelling to Bern with two kids aged 6 and 9 over Easter. We need two rooms that connect or adjoin, and somewhere we can walk to the bear park and the old town without a car.”

CONDITIONS IT CARRIES

walkable, no car

old town

connecting rooms

family of four in two rooms

WHAT WE ASKED • HOTEL IN BERN FOR A CONFERENCE OR EVENT

“I’m speaking at a conference in Bern in May and need somewhere I can also host a dinner for eight clients. Four star, walkable from the venue, with a decent restaurant and a room I can work in.”

CONDITIONS IT CARRIES

four star

central, walkable

private dining for eight

restaurant open when needed

desk, work space

Read the hollow tags. You have most of those qualities and your guests write about them all the time, yet none of them is stated anywhere a machine can read, which is why a longer and more realistic question would score you worse than this report does rather than better. That is the strongest argument we could make for doing this work, and we would rather you heard it as a limit of our own measurement than as a sales line.

09 • THE NEXT STEP

Nothing in this report needs you to buy anything. It is yours, and if you hand it to your own agency on Monday that is a fine outcome.

If you do want to go further, the useful next hour is a conversation rather than a proposal, covering who your best guest actually is, what you charge them by segment, and which of these ten questions you would throw out and replace. That hour is what turns four measured engines on a single Tuesday into something worth acting on.

And whoever looks after your marketing today, there is one question worth putting to them first, which is what your hotel's answer looked like in these four systems last Tuesday and what it looks like this Tuesday, because if they cannot tell you then they are not measuring it, and nobody fixes what they are not measuring.

10 • METHOD

How this was measured, and what we would not claim.

PERPLEXITY

60%

GEMINI

47%

CLAUDE

41%

OPENAI

37%

Thirty questions, ten intents in English, German and French, each phrased the way a native speaker types rather than translated. Asked on 30 September 2026.

Four engines through their own grounded search APIs: gpt-5-mini, perplexity/sonar, gemini-3.7-flash and claude-sonnet-5. 120 calls, of which 116 returned sources and were kept.

An answer that came back citing nothing was discarded rather than counted as an answer naming nobody. Three were discarded and one call failed. Two of the three were the same engine on German questions.

Per engine you are named: Perplexity 60%, Gemini 47%, Claude 41%, OpenAI 37%. The engine most people mean when they say AI search is the one that names you least.

How a mention was counted. An answer counts as naming you if it contains your name. That is reliable for one distinctive name and it is not a way to rank a market: these answers contain about sixty distinct hotel names, many of them the same hotel written two ways, and resolving those to a real league table is a separate pass we have not run.

Not yet measured: Google AI Overviews has no API and DeepSeek's search only works in its consumer app, so both are run by hand and are not in these figures. Every number here comes from a grounded API, which is a faithful proxy for the consumer product and is not identical to it.

No location was given to any engine, so a German question measures a German question from an unspecified place. Two of the four accept a location hint and two do not, and giving it to half of them would make the disagreement partly our own doing.

Every figure in this report traces to a stored answer with a date. Nothing here is an estimate except the three rates in the revenue table, which are labelled as such.

Three ways to get named.

We ask the engines about you every week. We fix what makes them skip you.
Each step up does a little more.

Presence

For one small hotel.

- We check if AI names you, every week.
- We fix what makes AI skip you.
- You get one page each month that shows how it went.
- You have one person you can call.

Direct Demand

MOST CHOSEN

For hotels guests travel to see.

- Everything in Presence.
- AI learns about you in English, French and German.
- We get you talked about on the sites AI reads.
- More guests book with you directly.

Signature

For luxury hotels and small groups.

- Everything in Direct Demand.
- A senior expert leads your work.
- We do more each month to get you named.
- We meet your owners every three months.

ALSO, WHEN YOU NEED IT

Once AI names you, we can help with four more things.

Branding

Your hotel looks and sounds the same everywhere.

Social media

We write posts that use the words guests search for.

First in Google Maps

Your hotel shows up at the top of the map, with the right photos and hours.

Workflow automation

Jobs you do by hand each week get done for you, like answering reviews fast.

Want to talk it through? We start with one hour.

Talk to us · [y8y.ai](#)