

MEASUREMENT

# You Can't See Whether AI Recommends You — And That's the Real Problem

By Scott Tischler · September 8, 2026 · 9 min read

---

**Here's a question most businesses cannot answer: when someone asks ChatGPT, Gemini, or Perplexity to recommend a business in your category, do you get named?** You probably track your Google rankings to the decimal. You almost certainly have no systematic idea what the AI answers say about you. That blind spot — not any single tactic — is the biggest problem in AI visibility right now, because you can't improve what you can't see.

The reason this matters more every month is that the traffic is genuinely moving. Gartner has projected a sharp decline in traditional search volume this year as people shift to AI-powered answers — on the order of a quarter of it. Whether the exact figure holds, the direction is not in dispute: a growing share of the moment where a customer decides who to consider now happens inside an AI answer you never see. If that's where decisions are forming and you're not measuring it, you're flying blind over the exact terrain that's starting to matter most.

## Why the old metrics don't cover this

Traditional SEO gave you a scoreboard. Rankings, impressions, clicks, positions — all observable, all trackable, all sitting in a dashboard. AI answers broke that scoreboard in three ways, and it's worth being precise about them.

First, **answers are personalized and non-deterministic**. Two people asking the same question can get different responses, and the same person asking twice can get different phrasing. There's no single "position 3" to occupy. Second, **most of it is invisible to your analytics**. When an AI answer mentions you but the user doesn't click through, nothing shows up in your traffic reports — you were recommended and you'll never know it from your logs. Third, **the surface is fragmented**. ChatGPT, Perplexity, Gemini, Google's AI answers, and others each behave differently, so there's no one place to look even if you wanted to.

The result is a strange situation: the channel that's growing fastest is the one you have the least visibility into. Your reporting is thorough on the surface that's shrinking and nearly silent on the one that's expanding.

## What "AI visibility measurement" actually means

Closing the gap doesn't require a crystal ball. It requires deliberately, repeatedly asking the questions your customers ask and recording what comes back. That sounds almost too simple, but the discipline of doing it systematically is what turns anecdote into signal. Concretely, measurement means tracking a few things over time:

- **Presence:** for the prompts that matter in your category, are you named at all? This is the binary floor — mentioned or not.
- **Share of answer:** when the model lists options, how often are you among them, and where? Being one of the few named sources is the whole game now.
- **Sentiment and framing:** when you are mentioned, how are you described? Being named as the budget option versus the premium expert is a real difference the raw "mention" count hides.
- **Accuracy:** is what the model says about you actually correct? Models sometimes confidently state wrong facts about a business, and you can't fix what you don't know is broken.
- **Competitors:** who gets named when you don't? Your competitors' presence in answers is as informative as your own.
- **Movement:** all of the above tracked over time, so you can tell whether what you're doing is working.

That last point is the one that turns measurement from a curiosity into a tool. A single snapshot tells you where you stand today. A trend line tells you whether your effort is paying off — which is the only thing that lets you invest confidently instead of guessing.

## The prompt set is the foundation

If you do one thing, build a prompt set: the actual questions a real customer would type when they're trying to find or choose a business like yours. Not the questions you wish they'd ask — the ones they actually ask. "Best [category] in [city]." "Who should I hire to [job]." "Is [your business] any good." "[Competitor] alternatives." A few dozen of these, chosen to reflect real buying intent, become your measurement instrument. Everything else is just running that instrument on a schedule and writing down the results.

The prompt set is also where honesty matters most. It's tempting to stack it with softball prompts where you already win. Resist that — a measurement system you've rigged to flatter you is worse than none,

because it gives you false confidence. Include the hard prompts, the competitor comparisons, the ones you're afraid of. Those are the ones with something to teach you.

## What you do with the signal

Measurement isn't the goal — it's what makes everything else improvable. Once you can see the answers, patterns emerge fast. Maybe you're absent from a whole cluster of prompts, which tells you where to build content and earn corroboration. Maybe you're mentioned but described inaccurately, which points at inconsistent facts across the web you need to reconcile. Maybe a competitor dominates a category you assumed you owned, which is a strategic wake-up call worth more than any ranking report. None of those insights are available to a business that isn't looking.

And there's a compounding benefit: once you're measuring, you stop arguing about AI visibility in the abstract and start managing it like any other channel — with a baseline, interventions, and a way to tell whether they worked. That shift, from vibes to instrumentation, is what separates the businesses that will quietly win this era from the ones still optimizing the channel that's shrinking.

## Start this week

You don't need a platform to begin. Write down the twenty questions your customers actually ask. Run them through the major AI answer engines yourself and record, honestly, what each one says — whether you're named, who else is, and how you're described. Do it again in two weeks. That crude baseline will already tell you more about your real market position than most of your existing dashboards, because it's measuring the place where decisions are increasingly made. From there you can decide what to fix — but at least you'll be able to see.

### Key takeaways

- ♦ Most businesses track Google rankings obsessively but have no systematic idea whether AI answers name them — that blind spot is the biggest problem in AI visibility.
- ♦ The traffic is genuinely moving: Gartner projects a sharp drop in traditional search this year as decisions shift into AI answers you never see.
- ♦ Old metrics don't cover this because AI answers are personalized and non-deterministic, mostly invisible to analytics, and fragmented across platforms.
- ♦ Measurement means tracking presence, share of answer, sentiment, accuracy, competitors, and movement over time — the trend line is what makes effort improvable.

- ♦ Build a prompt set of the real questions customers ask, and include the hard and competitor prompts — a measurement system rigged to flatter you is worse than none.
- ♦ You don't need a platform to start: run your top 20 customer questions through the major engines, record honestly what they say, and repeat in two weeks.

## Frequently asked questions

### Why can't I just use my website analytics to see AI visibility?

Because most AI mentions never generate a click. When an answer engine recommends you but the user doesn't visit your site, nothing appears in your traffic reports — you were recommended and your analytics stay silent. AI answers are also personalized and non-deterministic, so there's no fixed 'position' to track. You have to measure the answers directly, not infer them from downstream traffic.

### What's the single most important thing to set up first?

A prompt set: the actual questions a real customer would type when trying to find or choose a business like yours — 'best [category] in [city]', 'who should I hire to [job]', '[competitor] alternatives,' and so on. A few dozen intent-driven prompts become your measurement instrument. Run them through the major engines on a schedule and record what comes back.

### How often should I measure, and what should I record?

Start with a baseline now and re-run it every couple of weeks so you can see movement. For each prompt, record whether you're named at all, who else is named, where you fall in the list, how you're described (sentiment and framing), and whether what's said about you is accurate. Tracking those over time is what tells you if your efforts are working.

**About the author.** Scott Tischler is the Founder & Chairman of Alrecommend.ai and a practitioner-authority on AI search and Answer Engine Optimization. With 20+ years in marketing technology — including American Express, MetLife, and UBS — and executive and professional study at Wharton, Harvard, and Oxford, he helps businesses become the ones AI recommends.