



Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

AI & SEARCH

Your Competitor Is Being Recommended by AI Right Now — and You Have No Idea

By Scott Tischler · July 28, 2026 · 9 min read

The lead you'll never know you lost

Here's a scene playing out thousands of times a day, and you have no visibility into a single instance of it.

A buyer with a real budget and real urgency opens ChatGPT, Perplexity, Google's AI Overviews, or Gemini. They don't type keywords the way they used to. They type a full sentence: ***"Who's the best [what you do] for a mid-size company in [your city]?"*** The AI thinks for a second and answers in a calm, confident paragraph. It names two or three companies. It explains why they're a good fit. The buyer clicks one, or just reaches out directly.

If that company is your competitor, you didn't lose that deal in a pitch. You lost it before you knew it existed. There was no form fill, no bounce, no abandoned cart in your analytics. The customer was recommended away from you by a machine, in a conversation you were never part of.

I've spent more than twenty years in marketing technology — at American Express, MetLife, and UBS, and now running Alrecommend.ai — and I want to be blunt about what I'm seeing. This is the most under-measured competitive threat most businesses face right now. Not because it's subtle. Because it's ***invisible***. The whole category of loss doesn't show up in the dashboards executives are trained to watch.

Let me show you how to make it visible, why the gap exists, and how to close it — starting today.

This is not "SEO for chatbots"

The first thing I have to knock down is the reflex I hear in almost every intro call: ***"So this is just SEO for AI, right? We already do SEO."***

No. And the difference matters enormously.

Traditional search returns a **list**. Ten blue links, a menu of options, and the user does the comparing. Your job was to rank — to be one of the choices on the shelf. Even at position six, you were in the running.

AI search returns a **verdict**. It doesn't hand the user a shelf; it hands them a recommendation. It has already done the comparing, the filtering, and the summarizing. There is no position six. There's the answer the model gives, and then there's everyone it didn't mention — which is a silent, undifferentiated void. Being "on page one" of an answer that names one company means nothing if you're not that company.

This discipline has a name — Answer Engine Optimization, or AEO — and it's a genuinely different game. In Alrecommend.ai's **State of AI Search 2026** research, we found that **37% of buyers now begin their research with an AI tool rather than a traditional search engine**. More than a third of the top of your funnel is already flowing through a channel most companies have never audited even once.

And here's the part that should raise the hair on your neck: **between 58% and 68% of those AI searches end without a single click to any website**. The user gets their answer inside the AI and moves on. Zero-click. If the model named your competitor and paraphrased their expertise back to the buyer, your rival just got the benefit of a recommendation and you got nothing — not even a log entry. The traffic you'd normally use to **detect** that you're losing never arrives.

That's why this is invisible. You can't see the deals you're losing, because losing them doesn't generate any of the signals you're used to watching.

Why the gap exists (it's usually not your product)

When a company discovers a competitor is being consistently recommended over them, the first assumption is always the painful one: **they must be better**. Sometimes that's true. Usually it isn't. In my experience the gap is almost never about product quality — it's about **legibility to machines**.

Large language models don't experience your brand. They can't tour your office, feel your service, or be charmed by your founder. They assemble an answer out of the text that exists about you across the web, and out of how confidently and consistently that text describes what you do, who you're for, and why you're credible. A model recommends the company it can most easily and confidently **describe**.

So the competitor winning in AI answers is usually the one who is, in machine terms, easier to read. There are a handful of recurring reasons the gap opens:

- **They've published specific, structured expertise.** Not vague "we're passionate about quality" copy, but concrete answers to the exact questions buyers ask — with numbers, definitions, and

named use cases. Models feed on specificity.

- **They're described consistently everywhere.** Their category, their niche, and their differentiators read the same on their site, their directory listings, their reviews, and third-party mentions. Consistency reads as confidence to a model. Contradiction gets you dropped.
- **They have corroboration the model trusts.** Third parties — reviews, roundups, press, credible mentions — say similar things about them. An LLM weights claims that independent sources back up. Your "About" page alone is not corroboration.
- **Their content answers questions, not just describes services.** A page titled "Best options for X in [city]" that genuinely compares tradeoffs is far more citable than a brochure page titled "Our Services."

Notice none of these is "have a better product." They're all about whether the web has a clear, consistent, corroborated story about you that a machine can lift with confidence. Your competitor didn't necessarily out-build you. They out-**explained** you to the machines — probably by accident.

Run this audit today

You do not need to hire anyone or buy anything to find out where you stand. You need about forty-five minutes, a notepad, and a willingness to read answers that may sting. Do this now, before you read the fixes.

1. Build your buyer's question list. Write down the 10–15 questions a real prospect would ask an AI when they have your problem but don't yet know your name. Not **"[Your company] reviews"** — that's a branded search, and it's cheating. Use unbranded, intent-rich prompts: **"Best [category] for [customer type] in [location],"** **"How do I choose a [what you do],"** **"[Competitor] alternatives,"** **"Who should I hire to [job to be done]."**

2. Ask across at least three engines. Run every question through ChatGPT, Perplexity, and Google's AI Overviews at minimum; add Gemini and Copilot if you can. They pull from different sources and will not agree — that disagreement is data. Use a fresh session or logged-out/incognito window so your own history doesn't bias the answer toward you.

3. Record three things for every answer. For each prompt on each engine, note: **(a)** Were you named at all? **(b)** Which competitors were named, and in what order? **(c) How** were the named companies described — what specific claim or strength did the AI attribute to them? That third column is the goldmine. It tells you exactly which of your rival's characteristics the model found legible and worth repeating.

4. Follow the citations. In Perplexity and AI Overviews you can see the sources behind the answer. Open them. These pages are, quite literally, the content winning the AI's trust in your category. If your competitor's blog post, a directory, or a review site keeps appearing, you've just been handed the map of where the game is being played.

5. Score yourself honestly. Tally your "named" rate across all prompts and engines. If you're named in fewer than half of your own core buying questions, you don't have a minor visibility problem — you have a competitor being recommended in your place, repeatedly, to people you'll never get a second chance at.

When clients run this for the first time, the room goes quiet. Founders who are certain they "dominate" their niche watch an AI recommend three rivals by name and never mention them across a dozen prompts. That silence is the sound of the problem becoming visible for the first time. (If you'd rather see it tracked continuously across engines and prompts instead of by hand, monitoring your AI visibility is exactly what we built Alrecommend.ai to do — but do the manual pass first. Feeling it yourself changes how seriously you take the fix.)

How to close the gap

Once you can see where you stand, the work is concrete. Here's where I'd start, in order of leverage.

Answer the exact questions, in the exact structure, that models reward. Take the prompt list from your audit and build content that answers each one directly and specifically — clear question as the heading, a crisp answer in the first two sentences, then the supporting detail. Include real numbers, defined terms, named customer types, and honest tradeoffs. You're not writing to charm a human skimmer; you're writing to be the cleanest, most quotable source a machine can lift.

Make your description of yourself ruthlessly consistent. Audit how your category and differentiators are stated across your homepage, your listings, your profiles, and anywhere third parties describe you. Pick the precise language for who you serve and what you're best at, and enforce it everywhere. Every contradiction is a reason for a model to lower its confidence in recommending you.

Add the machine-readable layer. Implement structured data — schema for your organization, products, FAQs, and reviews. This isn't a ranking trick; it's translating your story into a format that removes ambiguity about what you are and who you're for. It lowers the cost for a model to describe you correctly.

Earn corroboration. Get credible third parties saying consistent things about you: reviews with substance, inclusion in the roundups and directories your audit surfaced, genuine expert commentary and press. The pages already cited in AI answers about your category are your target list. This is

slower work, but it's the difference between a model that **can** mention you and one that's **confident** enough to recommend you.

I want to be careful here, because this space is full of people promising overnight domination. I won't. What I'll offer instead is an informed prediction from the data we're seeing across engagements: companies that treat AEO as a full, sustained discipline — content, consistency, structure, and corroboration together — are seeing roughly **3.8× more citations in AI answers** and about **142% more AI-referred leads within 90 days**, per our **State of AI Search 2026** research. I frame that as directional, not a guarantee, because your category and starting point matter. But the direction is not in doubt.

The window is open, and it won't stay that way

Here's my honest read of the moment. Most of your competitors haven't run the audit above. The category is still soft. The businesses that move now — that make themselves legible while everyone else is still arguing about whether AI search is real — are setting the defaults the models will lean on for years. Recommendation is sticky; the company an AI learns to name confidently tends to keep getting named.

In two years, "we show up in AI answers" will be table stakes, fought over inch by inch. Right now it's available to whoever bothers to claim it.

So run the audit today. Not this quarter — today, in the next hour, with your real buyer questions. Best case, you're already the name the machines trust, and you'll defend it deliberately instead of by luck. Worst case, you'll watch an AI recommend your competitor to your customer, in your own words, right in front of you.

Either way, you'll finally be able to see the game you're already playing.

Key takeaways

- ② AI search returns a verdict, not a list — there's the company that gets recommended and the invisible everyone else. There is no "position six."
- ② 37% of buyers now start research with AI, and 58–68% of those searches end with zero clicks, so the deals you lose here never appear in your analytics.
- ② The gap is rarely about product quality; it's about being legible to machines — specific, consistent, structured, and corroborated across the web.

- ② Run the audit today: list 10–15 unbranded buyer questions, ask them across ChatGPT, Perplexity, and AI Overviews, and record who gets named and how they're described.
- ② Close the gap by answering exact buyer questions directly, enforcing consistent self-description, adding structured data, and earning third-party corroboration.
- ② The window is open now because most competitors haven't looked — recommendation is sticky, and early defaults are hard to dislodge later.

Frequently asked questions

How is Answer Engine Optimization (AEO) different from the SEO I already pay for?

SEO optimizes to rank on a list of links where the user does the comparing; AEO optimizes to be the company an AI names when it delivers a single recommendation. Traditional SEO helps but isn't sufficient — models reward specific, consistently described, corroborated, question-answering content, and they hand the buyer a verdict rather than a menu.

Why can't I see the leads I'm losing to AI recommendations?

Because losing them generates none of the signals you monitor. When an AI answers inside the chat and names a competitor, there's no click, no bounce, and no form fill on your site — 58–68% of AI searches are zero-click. The loss happens in a conversation you were never part of, which is why you have to audit the answers directly.

How long does it take to start showing up in AI answers?

It varies by category and starting point. Content and structured data can improve in days to weeks, while corroboration from third parties is slower. In Alrecommend.ai's State of AI Search 2026 research, companies running a full AEO program saw roughly 3.8x more citations and about 142% more AI-referred leads within 90 days. Treat that as directional, not a guarantee.

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 study at Wharton, Harvard, Yale, and Oxford, he helps businesses become the ones AI recommends.