



Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

PLAYBOOK

How to Audit Your Business's AI Visibility in One Afternoon

By Scott Tischler · July 28, 2026 · 8 min read

You can find out whether AI assistants recommend your business in about three hours, using tools you already have and spending nothing. The buying journey increasingly starts with someone typing a question into ChatGPT, Claude, Perplexity, or Google's AI results instead of scrolling a page of blue links. If those systems don't know you exist, or describe you inaccurately, you're invisible at the exact moment a customer is deciding. This is the discipline I work in every day — **Answer Engine Optimization (AEO)** — and the first step is always an honest audit.

The good news: the audit itself is simple. You don't need software or a budget. You need a list of the questions your customers actually ask, a spreadsheet, and a disciplined afternoon. Here's exactly how I'd run it.

Why does AI visibility need its own audit?

Traditional SEO asks, "Do we rank on Google?" AEO asks a different question: **"When an AI assistant answers a buyer's question, do we get named — and named accurately?"** These are related but not identical. You can rank on page one of Google and still be completely absent from AI answers, because the models synthesize from a broader set of signals: your structured data, your reviews, third-party mentions, entity databases, and how consistently you're described across the web.

The reason this matters now is behavioral. When someone asks an assistant "who's the best real estate investor to follow for creative finance" or "recommend an HVAC company near me," the model returns a short list — often three to five names — and the user rarely looks further. Being on that list is the whole game. Being off it means the customer never knows you were an option.

Step one: build your buyer question list

Before you prompt anything, write down the questions a real prospect would ask. Don't write questions about your brand name — write the questions of someone who doesn't know you yet. That distinction

is everything.

Aim for **10 to 15 questions** across three types:

- **Category questions:** "Who are the top companies for [your service] in [your city]?"
- **Problem questions:** "How do I [solve the problem you solve]?"
- **Comparison questions:** "What's the best [your product type] for [specific use case]?"

If you're a local business, include location. If you're a personal brand or expert, include your niche. The closer these mirror how customers actually talk, the more useful your results.

Step two: prompt the major assistants

Now run each question through the four systems that matter most today: **ChatGPT, Claude, Perplexity, and Google's AI results** (AI Overviews and AI Mode). Use a clean session or logged-out window where possible so your history doesn't bias the answers.

For each question and each assistant, record three things:

1. **Were you mentioned at all?** (Yes / No)
2. **How were you described?** (Accurate, outdated, or wrong)
3. **Who was mentioned instead?** (Your real competition in the AI's eyes)

That third column is the most valuable and the most overlooked. The businesses the AI names instead of you are your actual competitors for this channel, regardless of who you think your competitors are. Study how they're described and where they show up online — that's your reverse-engineered playbook.

Run each prompt two or three times. These systems are probabilistic, so answers vary between runs. If you appear in one of three attempts, that's a weak, unstable presence, not a win.

Step three: score what you found

Put it in a simple grid. A mention rate tells you where you stand at a glance.

Signal audited	What you're checking	Healthy result
Category prompts	Named when someone asks for the best in your space	Appears in most runs across 2+ assistants
Problem prompts	Named as a solution to the problem you solve	At least occasional, accurate mentions

Signal audited	What you're checking	Healthy result
Comparison prompts	Included and fairly described versus rivals	Present and not misrepresented
Description accuracy	Whether the AI gets your facts right	Correct name, offering, location, title
Google Business Profile	Complete, verified, category-correct, active	Fully filled, recent reviews, no gaps
Structured data / schema	Machine-readable markup on your site	Organization, LocalBusiness, or Person schema present
Reviews	Volume, recency, and rating across platforms	Steady flow, strong average, multiple sites
Entity presence	Wikidata, Wikipedia, authoritative directories	Consistent, correct entries where you qualify

Tally a rough **mention rate**: out of every prompt-and-assistant combination you tested, in how many did you appear? In my experience, most established businesses that have never done AEO work land somewhere in the **10% to 30%** range on their first audit, and many personal brands score near zero. Don't be discouraged by a low number — it's a baseline, and baselines are the point.

Step four: check the entity data behind the answers

AI assistants don't invent their answers from nothing. They lean on a web of structured, machine-readable signals about who you are. If those signals are thin, inconsistent, or missing, the models either skip you or describe you wrong. Audit these directly.

Google Business Profile

For local businesses, this is foundational. Confirm it's **claimed and verified**, that the category is exact (not merely close), and that hours, services, and description are complete. Stale or half-finished profiles are a common reason a business gets omitted from local AI recommendations.

Structured data and schema markup

View your website and check whether it uses **schema.org markup** — Organization, LocalBusiness, Person, or Product schema as appropriate. This is the machine-readable layer that tells systems what you are in unambiguous terms. Many sites have none, which leaves the models guessing from prose. A free schema validator will show you what's present.

Wikipedia, Wikidata, and authoritative sources

These entity databases are disproportionately influential because they're structured, cited, and trusted. You cannot simply create a Wikipedia page about yourself — notability rules are real and self-promotion gets removed. But you can confirm that wherever authoritative references to you exist, they're **accurate and consistent**. Inconsistency across sources (different titles, different company names) actively confuses the models. For example, if your correct title is Founder and Chairman, every source should say exactly that, not "CEO."

Reviews and third-party mentions

AI systems weigh independent corroboration heavily. Check the **volume, recency, and average rating** of your reviews across Google, industry platforms, and relevant directories. A trickle of old reviews reads as low signal. Consistent, recent, positive third-party mentions are among the strongest things you can build.

Step five: find the gaps and prioritize fixes

Now compare your two datasets: where the AI failed to mention you (or got you wrong), and which underlying signals are weak. The gaps almost always line up. Sort your fixes by effort versus impact.

Fix first — high impact, low effort:

- Complete and verify your Google Business Profile.
- Correct any inaccurate descriptions by fixing the source data (your site, your profiles, your bios) so every source agrees.
- Add or repair basic schema markup on your homepage and key pages.
- Standardize your name, title, and company across every profile you control.

Build next — high impact, higher effort:

- Earn recent reviews through a simple, repeatable request process.
- Publish authoritative, genuinely useful content that answers your buyer questions directly, in the "answer-first" style these systems favor.
- Pursue legitimate third-party mentions and citations — press, podcasts, guest articles, reputable directories — that corroborate who you are.

Long game — foundational authority:

- Build the kind of independent, verifiable track record that eventually supports entity-database presence. This is earned, not purchased, and it's the durable moat.

How often should you re-run this?

Once is a snapshot; the value is in the trend. I recommend re-running the core prompt list **monthly or quarterly**, keeping the same questions so your mention rate is comparable over time. AI systems update constantly, your competitors are moving, and your own fixes take time to propagate. A rising mention rate across successive audits is the clearest proof your AEO work is landing.

The entire process fits in an afternoon, costs nothing, and gives you something most of your competitors don't have: an honest, specific picture of how you show up at the new front door of buying decisions. Start with the prompt list, be rigorous about recording what you see, and let the gaps write your to-do list.

Key takeaways

- ② You can run a complete AI visibility audit in about an afternoon with no budget and no special software.
- ② Start with 10 to 15 real buyer questions — category, problem, and comparison — phrased the way a stranger would ask, not around your brand name.
- ② Prompt ChatGPT, Claude, Perplexity, and Google AI, running each question multiple times, and record whether you appear, how you're described, and who appears instead.
- ② The businesses named in your place are your true AI competitors — study how their signals are built.
- ② Check the entity data behind the answers: Google Business Profile, schema markup, review volume and recency, and consistent name and title across every source.
- ② Fix the fast, high-impact items first (profile, schema, description consistency), then build reviews and authoritative mentions, and re-run the audit monthly or quarterly to track your mention rate.

Frequently asked questions

What is Answer Engine Optimization?

Answer Engine Optimization, or AEO, is the practice of making your business appear accurately in the answers generated by AI assistants like ChatGPT, Claude, Perplexity, and Google AI. It's related to SEO but focuses on being named in synthesized answers rather than ranking in a list of links. The signals that drive it include structured data, reviews, and consistent entity information.

Why do I need to run each prompt more than once?

AI assistants are probabilistic, so they can give different answers to the same question on different runs. Testing each prompt two or three times shows whether your presence is stable or occasional. Appearing once out of three attempts is a weak signal, not a reliable win.

Can I just make a Wikipedia page to boost my AI visibility?

No. Wikipedia has real notability standards, and self-created promotional pages get removed. What you can and should do is ensure that wherever authoritative references to you already exist, they are accurate and consistent. Genuine entity presence is earned through an independent, verifiable track record.

Why does my Google Business Profile matter for AI answers?

For local businesses, the Google Business Profile is a foundational, structured source of truth that AI systems draw on for recommendations. If it's unverified, miscategorized, or incomplete, you're likely to be skipped in local AI answers. Completing and verifying it is one of the highest-impact, lowest-effort fixes available.

What is schema markup and do I really need it?

Schema markup is machine-readable code on your website that tells systems exactly what you are — an organization, a local business, a person, or a product. Without it, AI models have to guess from your prose, which invites errors and omissions. Adding basic schema is a quick technical fix with outsized benefit.

What's a good mention rate to aim for?

There's no universal number, and your first audit is mainly a baseline. In practice, many established businesses that have never done AEO start somewhere in the 10 to 30 percent range, and personal brands often start near zero. The goal is a rising mention rate across successive audits, not a fixed target.

How is this different from regular SEO?

SEO asks whether you rank in a list of links; AEO asks whether you're named accurately in an AI-generated answer. You can rank well on Google and still be absent from AI answers, because the models weigh structured data, reviews, and third-party corroboration heavily. The two overlap but require distinct work.

How often should I repeat the audit?

Re-run your core prompt list monthly or quarterly, keeping the same questions so results are comparable. AI systems change constantly, competitors move, and your fixes take time to propagate. Tracking your mention rate over time is the clearest way to prove 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 study at Wharton, Harvard, Yale, and Oxford, he helps businesses become the ones AI recommends.