



Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

AI & SEARCH

AI Is Rewriting Your Brand Right Now — Here's How to Take Back Control

By Scott Tischler · July 28, 2026 · 10 min read

There's a sales conversation happening about your brand at this exact moment, and you weren't invited.

A buyer opens ChatGPT, Gemini, Claude, or Perplexity and types something like "best [your category] for a mid-size company" or "is [your company] any good?" In a second and a half, the model produces a confident, fluent, authoritative-sounding paragraph about who you are, what you sell, how you compare, and whether you're worth the money. The buyer reads it as fact. And you never saw it, can't see the next thousand versions of it, and — until recently — had no real way to influence it.

I've spent more than twenty years in marketing technology, including inside Amex, MetLife, and UBS, and I've watched a lot of channels come and go. This one is different. AI assistants aren't just a new place to be found. They're a new **narrator** — one that summarizes, judges, and recommends on your behalf, to an audience you can't see, using information you didn't approve.

According to Alrecommend.ai's State of AI Search 2026 research, 37% of buyers now begin their product and vendor research inside an AI assistant rather than a traditional search box. That's not a fringe behavior anymore. And here's the part that should keep you up at night: the same research finds that between 58% and 68% of AI-assisted searches are "zero-click" — the buyer gets their answer from the model's summary and never visits a single website. They form an opinion about you without ever touching a page you control.

If the model gets you right, that's free, high-trust marketing. If it gets you wrong — outdated pricing, a product you sunset two years ago, a competitor's feature attributed to you, or a flat-out hallucinated "fact" — that's a silent, compounding liability. The good news: this is manageable. You can audit what AI says about you, correct what's wrong, and actively shape the description. Let me show you how.

How AI Actually Forms Its Picture of You

To take control, you have to understand what you're controlling. AI assistants build their description of your brand from three overlapping sources.

1. Training data (the model's long-term memory). Large language models are trained on enormous snapshots of the public web, frozen at a point in time. Whatever was widely written about you before that cutoff — your old website copy, press coverage, reviews, forum threads, directory listings — is baked into the model's parameters. This is why models sometimes describe your company as it existed 18 months ago. They're recalling, not looking. You can't edit training data directly, but you heavily influence the *next* training run by changing what the web says about you now.

2. Retrieval (the model's live research). Most modern assistants don't rely on memory alone. When a buyer asks about you, they run live searches, pull current pages, and summarize what they find in real time. This is retrieval-augmented generation, and it's your biggest lever, because it uses today's web, not last year's snapshot. If your current, accurate information is easy to find and machine-readable, retrieval will favor it.

3. Entity signals (how the machine knows you're you). Behind the scenes, AI systems try to resolve your brand to a distinct "entity" — a stable node with attributes: what you do, who you serve, where you're located, who founded you, what you're known for. These signals come from structured data on your site, knowledge bases like Wikidata and Wikipedia, business directories, and the consistency of how you're described across the web. Weak or contradictory entity signals are the single most common cause of AI getting your brand wrong. If three sources say three different things about your category or headquarters, the model picks one — and it might not pick the true one.

The practical takeaway: your AI narrative is the product of what the web remembers about you (training), what the web currently says (retrieval), and how confidently the machine can identify you (entity signals). Fix all three and you're not gambling anymore.

Step One: Audit What AI Says About You Right Now

You can't correct a narrative you haven't read. Before you change anything, spend an afternoon doing a structured AI brand audit. This is the exact method I recommend.

Pick your platforms. Run every test across at least four assistants: ChatGPT, Gemini, Claude, and Perplexity. They pull from different sources and will disagree — and those disagreements are diagnostic gold.

Run the four question types. For each platform, ask:

- **Direct identity:** "What is [company]? What do they do and who do they serve?"

- **Reputation:** "Is [company] reputable? What do customers say about them? What are the downsides?"
- **Comparison:** "How does [company] compare to [competitor A] and [competitor B]?"
- **Recommendation:** "What are the best [your category] options for [your ideal customer]?" — and note whether you appear at all.

Log the results in a simple grid. Columns for each platform, rows for each question. In each cell, capture three things: Is the core description accurate? What's wrong, missing, or stale? And critically — **what source did it cite?** Most assistants now show citations. Those cited URLs are the exact pages shaping your narrative, which tells you precisely where to focus.

Score the damage. Flag every cell as Green (accurate and on-message), Yellow (technically true but outdated or thin), or Red (wrong, hallucinated, or you're invisible in a recommendation query where you should appear). Your Reds are emergencies. Your Yellows are the opportunity. Green means the machine is already selling for you — protect it.

Watch for the invisible failure. The worst outcome usually isn't a lie. It's absence. If you ask "best options for [ideal customer]" and you're simply not mentioned while three competitors are, that's a zero-click buyer who never learned you exist. Given that most AI searches never produce a website visit, being left out of the answer is being left out of the market.

Do this quarterly at minimum. Model outputs drift, competitors move, and your audit is only a snapshot.

Step Two: Correct What's Wrong or Stale

Now fix the Reds and Yellows. The order matters — start where the machine actually looks.

Own your entity home base. Publish a single, unambiguous "source of truth" about your company: what you do, who you serve, your categories, founding, leadership, locations, and key facts. Put it on a canonical page and reinforce it with structured data (Organization and Product schema markup). This gives retrieval a clean, authoritative page to prefer and strengthens the entity signals that help the machine identify you correctly.

Fix the cited sources directly. Your audit told you which URLs the models quoted. If a model is citing a stale third-party directory listing that says you're still in a category you left, update the listing. If it's pulling an outdated product description from your own site, rewrite it. You're not arguing with the AI — you're editing its reading material.

Correct the knowledge graph. Wikidata is publicly editable and disproportionately influential on how AI systems resolve entities. Make sure your core facts there are accurate and sourced. If you have a

Wikipedia presence, ensure it's factual and current (following their rules — no self-promotional edits).

Resolve contradictions. Hunt down places where the web disagrees about you — old category descriptions, a former product name, an outdated tagline, a stale address — and reconcile them. Consistency is what lets the machine speak about you with confidence instead of hedging or guessing wrong.

Feed the retrieval layer fresh material. Publish current, substantive content that directly answers the questions buyers actually ask an assistant. When your accurate content is the most retrievable content, retrieval does your correcting for you on every future query.

Step Three: Actively Shape the Description

Correcting errors gets you to neutral. Shaping the narrative is how you win. This is what Answer Engine Optimization (AEO) really is — deliberately structuring your presence so AI assistants describe you accurately, favorably, and often.

The results justify the effort. Alrecommend.ai's State of AI Search 2026 research found that brands implementing a full AEO program earned 3.8× more AI citations and roughly 142% more AI-referred leads within 90 days. That's not from gaming a model. It's from making yourself the clearest, most credible, most machine-legible answer to the questions your buyers ask.

To shape the description:

Write the sentences you want repeated. Models summarize and paraphrase. Give them clean, quotable, factual statements about your positioning, differentiators, and ideal customer — in plain language, high on the page. Vague, adjective-stuffed copy gets discarded; specific, verifiable claims get repeated.

Build genuine, citable authority. AI assistants weight sources they trust. Earned media, expert bylines, original research, and third-party validation all raise the odds that the machine both cites you and speaks about you with confidence. This is real reputation work, not a hack.

Answer the comparison and recommendation queries head-on. These are the highest-intent questions and the ones that decide deals. Publish honest, substantive content that helps a buyer choose, including where you fit and who you're best for. Assistants favor sources that genuinely help the user decide.

Match structure to how machines read. Clear headings, direct question-and-answer formatting, structured data, and factual precision all make your content easier to retrieve and quote accurately.

The Common Mistakes

I see the same errors repeatedly, and they're all avoidable.

- **Assuming SEO covers it.** Ranking #1 on Google does not guarantee the AI mentions you. Different systems, different signals. Traditional SEO is necessary but not sufficient.
- **Auditing once and stopping.** This is a monitoring discipline, not a project. Outputs drift weekly. A one-time check is a one-time snapshot.
- **Trying to argue with the model instead of editing its sources.** You can't scold an AI into accuracy. You change what it reads.
- **Ignoring absence.** Teams obsess over the occasional wrong fact and completely miss that they never appear in recommendation queries at all. Invisibility costs more than error.
- **Chasing manipulation over substance.** Prompt-injection tricks and keyword-stuffing get filtered out and can damage trust signals. The durable strategy is being genuinely the best, clearest answer.
- **Treating it as marketing's problem alone.** Your AI narrative touches product accuracy, PR, legal, and support. It needs an owner with cross-functional reach.

The Window Is Open Now

A little over a third of your buyers already start with an AI assistant, and most of them never click through to hear your side. The machine is the first — and often the only — voice they hear about you. Right now, most brands are letting that voice improvise from whatever scraps of the web it happens to find.

You don't have to. Audit what the assistants say, correct the sources they read, and deliberately shape the description until the machine tells your story the way you'd tell it yourself. The brands that do this in the next year will own their AI narrative. The ones that wait will spend the following year trying to correct one the machine wrote for them.

Take the pen back.

Key takeaways

- ② AI assistants already narrate your brand to buyers, and 37% of buyers now start research inside an AI tool while 58–68% of those searches never produce a website click (Alrecommend.ai's State of AI Search 2026 research).
- ② Models describe you using three sources: frozen training data, live retrieval, and entity signals —

retrieval and entity signals are the levers you can move today.

- ② Audit quarterly across ChatGPT, Gemini, Claude, and Perplexity using four question types (identity, reputation, comparison, recommendation), and log every citation the model shows.
- ② Fix errors by editing the cited sources, publishing a canonical source-of-truth page with structured data, correcting Wikidata, and resolving contradictions across the web.
- ② The most expensive failure is usually absence, not error: if you're missing from recommendation answers, buyers never learn you exist.
- ② A full AEO program drove 3.8× more AI citations and roughly 142% more AI-referred leads in 90 days (Alrecommend.ai's State of AI Search 2026 research) — earned through substance and machine-legibility, not tricks.

Frequently asked questions

How is Answer Engine Optimization (AEO) different from SEO?

SEO optimizes to rank a link on a search results page; AEO optimizes to be accurately described, cited, and recommended inside an AI assistant's generated answer. They share fundamentals like authority and clean structure, but ranking well on Google does not guarantee an AI mentions you — different systems weigh different signals, and most AI answers never send the user to a page at all.

Can I force an AI to remove false information about my company?

Not directly, and arguing with the model won't work. AI assistants generate answers from their training data and live retrieval, so the reliable fix is to correct the underlying sources: update the specific URLs the model cites, publish an authoritative source-of-truth page, correct entries in public knowledge bases like Wikidata, and reconcile contradictory information across the web. Change what the machine reads and its answers follow.

How often should I audit what AI says about my brand?

At least quarterly, and more often if you're in a fast-moving category, have just rebranded, changed pricing, or launched or sunset a product. Model outputs drift, competitors adjust their own presence, and retrieval pulls from a constantly changing web, so any single audit is only a snapshot of that moment — treat it as ongoing monitoring, not a one-time project.

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.

