



Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

PLAYBOOK

How to Get Your Business Cited by ChatGPT, Claude, and Perplexity

By Scott Tischler · July 28, 2026 · 11 min read

The New Front Door Is a Sentence, Not a Page

Ten years ago I taught my teams to fight for position one on Google. Today I teach them a harder truth: the buyer may never see a search results page at all. They ask ChatGPT for a shortlist, read Claude's synthesis, or watch Perplexity assemble an answer with three or four cited sources — and they act on it.

In Alrecommend.ai's **State of AI Search 2026** research, we found that 37% of buyers now start a purchase or vendor-research journey with an AI assistant rather than a traditional search engine, and 58–68% of those AI sessions end without a single click to an outside website. That is the zero-click reality, and it changes the job. The old game was ranking. The new game is being **named** — cited in the answer, recommended in the shortlist, quoted as the source. That discipline is Answer Engine Optimization (AEO), and after building AEO programs across dozens of B2B categories, I can tell you it is learnable, measurable, and mostly unclaimed by your competitors.

Below is the exact system I use. No theory you can't act on this quarter.

First, Understand How the Machine Decides to Name You

You cannot optimize for a mechanism you don't understand, and the three systems in the title do not work the same way. Group them by how they source answers.

Retrieval / answer engines — Perplexity and Google AI Overviews. These run a live search, pull a handful of pages, and synthesize an answer with visible citations. The citation is the whole product. Your job here overlaps with SEO — you must rank in the underlying index — but it goes further: your page has to be the one whose passage is **easiest to lift** cleanly into the answer. Perplexity and AI Overviews reward a page that states the answer plainly, near the top, in language that can be quoted without editing.

Chat models — ChatGPT and Claude. Two modes matter. When these tools browse (ChatGPT with search, Claude with web access), they behave like retrieval engines and cite live sources — same rules apply. When they answer from training and memory without browsing, there is no live citation; the model recommends whatever it has **learned** about your category. You earn a place there by being mentioned so consistently, across so many credible sources, that your brand becomes part of what the model "knows." You cannot inject yourself into a training run. You can flood the open web with enough consistent, authoritative signal that the pattern forms on its own.

The strategic takeaway: retrieval rewards **page-level extractability**; training rewards **web-wide consensus**. A serious program works both. Everything that follows serves one or both of those goals.

Write So a Machine Can Lift Your Answer Cleanly

Most B2B content is built to impress a human who is already reading. AEO content is built to be **extracted** by a system deciding whether to quote you at all. Different craft.

Lead with the answer. For every question your buyer asks, put a direct, self-contained answer in the first 40–60 words under the relevant heading — before the context, story, or nuance. Answer engines lift the passage that resolves the query fastest. If your best sentence is in paragraph six, it loses to a competitor who put theirs first.

Make headings the actual questions. Use the literal phrasing your buyers use: "How much does X cost?" "What is the difference between A and B?" Then answer immediately underneath. This maps your content to the query and hands the model a clean question-answer pair.

Structure for lifting. Definition sentences ("X is a..."), comparison tables, numbered steps, and tight bulleted criteria get extracted far more often than dense prose. When I convert a wall-of-text page into clear Q&A blocks with one idea per paragraph, its citation rate climbs — same facts, more liftable form.

Be self-contained and specific. Avoid "as we discussed above" and vague antecedents; a model may lift a passage in isolation, so each one must stand alone. Include concrete numbers, dates, named methods, and units. Specificity reads as credibility to a system weighing which source to trust.

State your expertise on the page. Name the author, their credentials, and the firsthand basis for the claim — "in the 40+ AEO audits we ran in 2025..." Firsthand experience is the E in E-E-A-T, and it is exactly what distinguishes a citable source from generic filler.

Give the Machines Structured Data to Remove Doubt

Structured data (Schema.org markup) does not force a citation, but it removes ambiguity about what your page contains — and ambiguity is what makes a model skip you. Implement, at minimum:

- **Organization** and **sameAs** — your official name, logo, and links to your authoritative profiles (LinkedIn, Crunchbase, your Wikipedia entry if you have one), so systems can resolve "who is this company" to one confident entity.
- **FAQPage** — on pages built as question-answer blocks, so each pair is machine-legible.
- **Article** with author and datePublished — establishing authorship and freshness.
- **Product / Service / Review** — where you want the model to understand what you sell and how it's rated.

Then get the plumbing right: keep the page crawlable, keep answers in real HTML text (never locked inside images or client-side-only rendering), and check your `robots.txt` against the AI crawlers you actually want — GPTBot, ClaudeBot, PerplexityBot, Google-Extended. I have watched teams spend a quarter on brilliant content that no engine could read because one line in `robots.txt` quietly blocked the bot. Verify it.

Get Into the Sources These Tools Already Trust

Here is the hardest truth in this playbook: your own website is rarely enough. When AI systems answer category questions — "best X for mid-market," "top vendors in Y" — they lean heavily on third-party sources they trust: industry roundups, review platforms, reputable editorial, structured databases, and the community forums where your buyers actually talk.

So map the sources for your category, then earn presence in them:

- **Review and directory platforms** — G2, Capterra, Gartner Peer Insights, and the niche directories specific to your industry. These are dense with the comparative language AI engines quote when building a shortlist.
- **Editorial and trade press** — the publications that already rank for your category terms and get pulled into AI Overviews and Perplexity answers.
- **Community sources** — Reddit, Wikipedia, Stack Overflow, and specialist forums carry outsized weight because models treat them as candid, human consensus. You cannot fake your way in; you earn it by being genuinely present and useful.
- **Wikipedia and structured knowledge bases** — if your company legitimately meets notability standards, a well-sourced entry becomes a foundational trust signal many systems reference. Earn it honestly; do not manufacture it.

The goal is simple to state and hard to do: when an AI reads the ten sources it trusts for your category, your name should appear in a meaningful share of them.

Build Authority Off-Site: Digital PR and Consistent Mentions

This is where retrieval and training strategy converge. Every credible, on-topic mention of your brand across the open web does double duty — it can be cited live **and** it strengthens the pattern the models learn from. Treat earned authority as a core AEO channel, not a vanity exercise.

What actually moves the needle:

- **Original data and research.** Publish proprietary numbers — survey results, benchmarks, index reports — that others cite. A defensible statistic gets quoted and re-quoted, and each citation carries your name into more of the web the models read.
- **Expert commentary and bylines.** Get your named experts quoted in trade press and contributing signed articles on authoritative sites. This builds the person's entity **and** the company's.
- **Consistent entity signals.** Use the same company name, boilerplate, and category description everywhere. Inconsistency splits your identity across multiple weakly-defined entities; consistency compounds one strong one.
- **Podcasts, panels, and expert roundups.** Being named alongside recognized players in your category teaches the model that you belong in that set — which is precisely what you want when someone asks for a shortlist.

Do this cleanly. No fake reviews, no manufactured quotes, no paid mentions dressed as editorial. Beyond the ethics, models and platforms increasingly discount low-trust signals, so the manipulative route underperforms the honest one anyway.

Close the Loop: Measure, Diagnose, Repeat

You cannot manage what you cannot see, and AI citations do not show up in your normal analytics — the zero-click sessions never reach your site. Build a deliberate measurement loop.

1. **Track your citation share.** For your 20–50 highest-value buyer questions, query ChatGPT, Claude, Perplexity, and Google AI Overviews on a fixed cadence and log whether you're cited, recommended, or absent — and which competitors appear instead. This share-of-voice is your core KPI. (This monitoring is exactly what we built Alrecommend.ai to automate, but you can start manually in a spreadsheet this week.)
2. **Watch the AI referral trickle.** Of the sessions that do click, some reach you. Segment referral traffic from AI domains in your analytics and track its trajectory.
3. **Diagnose the misses.** When a competitor is cited and you're not, read the answer's actual sources. Is it a review site you're weak on? A comparison page you never wrote? A data point they own? The

gap tells you the next move.

4. **Ship against the gap, then re-measure.** Improve the page, earn the placement, publish the data — then re-run the query set to confirm the citation moved.

This loop is where the returns compound. In our **State of AI Search 2026** research, businesses running a full AEO program — extractable content, structured data, trusted-source presence, and this measurement loop together — saw 3.8× higher AI citation rates and roughly 142% more AI-referred leads within 90 days than those doing piecemeal work. The multiplier comes from doing the whole system, not one tactic.

Your 90-Day Action Plan

1. **Audit crawlability (week 1).** Confirm GPTBot, ClaudeBot, PerplexityBot, and Google-Extended can reach your key pages; fix `robots.txt` and rendering issues first.
2. **Build your question set (week 1).** List the 20–50 real questions your buyers ask before purchase. This drives everything downstream.
3. **Baseline your citation share (week 2).** Query all four systems against that set and record where you stand today.
4. **Rewrite your top 10 pages for extractability (weeks 2–5).** Question headings, answer-first passages, tables, self-contained paragraphs, visible author credentials.
5. **Deploy structured data (weeks 3–6).** Organization, FAQPage, Article, Product/Review across priority pages.
6. **Map and enter your trusted sources (weeks 4–10).** Fix your review-platform presence; earn placement in the editorial and community sources AI actually pulls.
7. **Launch one digital-PR asset (weeks 6–12).** Publish original data or research designed to be cited, and pitch named experts for quotes and bylines.
8. **Re-measure and iterate (week 12, then monthly).** Re-run the question set, compare citation share to baseline, and ship against the biggest remaining gap.

Common Mistakes I See Every Week

- **Treating AEO as SEO with a new name.** Ranking helps retrieval engines but does nothing for a non-browsing model's recommendations. You need the off-site authority layer too.
- **Burying the answer.** Beautiful long-form that makes the model dig. If the answer isn't liftable in the first 60 words under a clear heading, you lose the citation.

- **Blocking the crawlers by accident.** The single most common own-goal. Check it before anything else.
- **Ignoring third-party sources.** Pouring everything into your own site while the review platforms and forums the AI trusts never mention you.
- **Manufacturing authority.** Fake reviews and planted mentions are low-trust signals that get discounted — and they carry real reputational risk. Earn it.
- **Not measuring.** Without a citation-tracking loop you're guessing, and you'll never know which change actually earned the mention.

The window here is the opportunity. Most of your competitors are still optimizing for a search results page fewer buyers visit. Do the full system — extractable content, structured data, trusted-source presence, earned authority, and a real measurement loop — and you become the source the machine names while they wonder where their pipeline went.

Key takeaways

- ② Retrieval engines (Perplexity, Google AI Overviews, browsing ChatGPT/Claude) reward page-level extractability; non-browsing models reward web-wide consensus — a real program works both.
- ② Lead every page with a self-contained, 40–60-word answer under a question-phrased heading; that single move drives more citations than any other.
- ② Your own site is rarely enough — you must be present in the review platforms, editorial, and community sources these tools already trust.
- ② Structured data and clean crawlability remove doubt; an accidental `robots.txt` block is the most common own-goal.
- ② Treat digital PR and consistent brand mentions as a core AEO channel — earned honestly, never manufactured.
- ② Track citation share across all four systems on a fixed cadence, diagnose the misses, ship against the gap, and re-measure — the loop is where returns compound.

Frequently asked questions

How is Answer Engine Optimization different from SEO?

SEO aims to rank a page on a results page a user then clicks. AEO aims to get your business named directly in an AI answer, whether the user ever clicks or not. It requires SEO's crawlability and ranking fundamentals plus

page-level extractability and off-site authority so that non-browsing models recommend you too.

Can I pay to be cited by ChatGPT, Claude, or Perplexity?

No — there is no paid placement in their organic citations, and attempting to game it with fake reviews or planted mentions produces low-trust signals that get discounted. Citations are earned through extractable content, structured data, and genuine presence in the sources these systems already trust.

How long before I see results from an AEO program?

Page-level improvements can affect retrieval-based citations (Perplexity, AI Overviews) within weeks of re-crawling. Off-site authority and non-browsing model recommendations build more slowly. In Alrecommend.ai's research, full-program teams saw meaningful gains within 90 days — but only when they ran the whole system rather than one tactic.

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.