



Scott Tischler
FOUNDER & CHAIRMAN · AIRECOMMEND.AI

PLAYBOOK

The Exact Prompts Your Customers Are Typing Into AI About Your Industry

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

Your customers are asking AI assistants about your industry right now, in full sentences, and getting recommendations that may never mention you. That is the shift. Search stopped being a list of blue links you optimize for and became a conversation you either show up in or you don't.

I build systems that get companies cited by AI answer engines, and the single biggest gap I see is this: **companies still think in keywords while their buyers think in prompts.** This playbook closes that gap. It shows you how to find the exact prompts people type into AI about your category, map that space, and engineer the signals that make you the answer.

Prompts are not keywords — and the difference is the whole game

A keyword is a fragment. A prompt is a full request with context, intent, and constraints baked in.

Someone searching Google in 2018 typed `crm for small business`. That same person now types into an AI assistant: ***"I run a 12-person agency and we're outgrowing spreadsheets — what CRM should we use that's affordable and easy to set up?"*** Same underlying need, completely different surface.

	Keyword era	Prompt era
Input	2–4 word fragment	Full natural-language question with context
Output	10 blue links	One synthesized answer, sometimes with a shortlist
User intent	Inferred	Stated explicitly in the prompt
How you win	Rank on the page	Get cited in the answer
Real estate	Page 1 has 10 slots	The answer names 1–5 options

That last row is the one that should keep you up at night. **A search results page has ten spots. An AI answer often names three.** The competition for inclusion is brutal, and being "on page one" means nothing if the model never surfaces you.

Answer Engine Optimization (AEO) is the practice of engineering your content and entity signals so AI assistants cite you as the answer to the prompts your buyers actually ask. That is the job. Everything below is how you do it.

Step 1: Discover the real prompts

You cannot win a prompt space you have not mapped. Start by building the actual list of questions buyers ask AI about your category. Real sources, in order of value:

- **Talk to your customers and sales team.** The questions prospects ask on discovery calls are, almost word for word, the prompts they type into AI later. Your sales notes are a goldmine.
- **Mine support tickets, chat logs, and your site search.** These are unfiltered, real-language questions from real buyers.
- **Ask the assistants directly.** Prompt ChatGPT, Claude, Gemini, and Perplexity with *"What are the most common questions people ask when choosing a [your category]?"* and *"What do buyers want to know before purchasing [your product type]?"* The models will hand you a starter map of the space.
- **Read the "People also ask" and community sources** — Reddit, industry forums, review sites. This is where buyers ask each other the questions they later ask AI.

Now expand each seed into the full prompt, with context and constraints. Don't stop at *"best project management tool."* Capture the real shape: *"best project management tool for a remote design team under \$15 a user."* The context is where the buying intent lives.

Step 2: Map the prompt space

Once you have a list, organize it. A raw pile of prompts is not a strategy; a mapped space is. I group prompts along two axes: **buyer intent** and **prompt type**.

By intent stage:

- **Problem-aware prompts:** *"Why is my team missing deadlines?"* — the buyer doesn't yet know your category is the answer.
- **Solution-aware prompts:** *"What kind of software helps teams hit deadlines?"* — they know the category, not the vendor.

- **Vendor-comparison prompts:** *"Is [Competitor A] or [Competitor B] better for agencies?"* — they are shortlisting, and this is where citations convert.
- **Validation prompts:** *"Is [Your Brand] any good?"* — they are checking you out. What the model says here can make or break the deal.

By type: **recommendation prompts** ("what's the best..."), **comparison prompts** ("X vs Y"), **how-to prompts**, **definition prompts**, and **fit prompts** ("best X for [specific situation]").

Prioritize the prompts that sit closest to a purchase decision and where you have a genuine right to win. A defensible comparison prompt in your niche is worth more than a broad recommendation prompt you'll lose to giants. Score each prompt on buyer value and your credibility to answer it, and work the high-value, winnable ones first.

Step 3: Become the answer

Here is the core mechanic: **AI assistants assemble answers from content they can parse and entities they trust.** So you win on two fronts — content that answers the prompt cleanly, and entity signals that tell the model you are a credible source.

On the content side:

- **Answer the prompt directly, up front.** Lead with the answer in the first sentence, then support it. Models lift self-contained, quotable statements. Buried answers get skipped.
- **Structure for extraction.** Clear question-shaped headings, short paragraphs, definitions, comparison tables, and lists. Make it trivial for a model to pull a clean chunk.
- **Cover the whole prompt space, not one page.** Build genuinely useful content for each priority prompt — comparisons, fit guides, definitions, honest "who this is not for" sections.
- **Be specific and honest.** Real numbers, real constraints, real tradeoffs. Models and buyers both reward content that reads like it was written by someone who actually knows the domain.

On the entity side — this is what most people miss:

- **Build a consistent, machine-readable identity.** Your brand, what you do, and who you serve should be stated the same way everywhere: your site, your About page, structured data, and your profiles.
- **Earn third-party mentions.** Models weigh what **others** say about you — reviews, mentions on credible sites, being listed in roundups and comparisons. Independent corroboration is the strongest entity signal there is.

- **Get the facts about you right across the web.** Consistent descriptions across review platforms, directories, and reputable publications teach the model who you are and what you're good at.

The blunt version: **content makes you eligible to be cited; entity signals make you trusted enough to be chosen.** You need both.

Step 4: Monitor your share-of-model

You cannot manage what you don't measure, and in AEO the metric is not rankings — it is **share-of-model: how often, and how favorably, AI assistants mention and recommend you across your priority prompts.**

Build a simple tracking loop:

1. **Take your priority prompt list** from Step 2 — the ones that matter most.
2. **Run them across the major assistants** on a regular cadence — monthly at minimum, weekly for the highest-value prompts.
3. **Record three things for each:** Are you mentioned? Are you recommended (not just named)? What exactly does the model say about you, and who does it cite?
4. **Track the trend and the gaps.** Where are competitors named and you're not? What prompts do you lose, and why? Which corrections move the needle?

This is a repeatable measurement discipline, not a one-time audit. Answers shift as models update and as the web around you changes, so **treat share-of-model like a dashboard you watch, not a report you file.** When you see a prompt where a competitor is winning, trace it back: is it a content gap, an entity-signal gap, or a corroboration gap? Then fix that specific thing and watch the next run.

Put it together

The companies that win the next decade of search will not be the ones with the most keywords. They will be the ones who **mapped the prompt space, built genuinely useful content against it, earned the entity signals that make models trust them, and measured their share-of-model relentlessly.**

Your buyers are already asking. The only question is whether the answer includes you. Start with ten real prompts this week — pull them from your last ten sales calls — run them through the assistants, and see what the machines say about you right now. That gap between what they say and what you want them to say is your entire roadmap.

Key takeaways

- ② Buyers now ask AI full, context-rich prompts, not keyword fragments — and AI answers name a few options, not ten links.
- ② Answer Engine Optimization means engineering content and entity signals so AI cites you as the answer to real buyer prompts.
- ② Discover prompts from sales calls, support tickets, community forums, and by asking the assistants directly.
- ② Map the prompt space by buyer intent and prompt type, then prioritize high-value prompts you can genuinely win.
- ② Win on two fronts: content that answers the prompt cleanly and entity signals that make models trust you.
- ② Track share-of-model — how often and how favorably AI names and recommends you — on a regular cadence.

Frequently asked questions

What is the difference between a keyword and a prompt?

A keyword is a short fragment like "crm software," while a prompt is a full natural-language question with context and constraints, like "what CRM should a 12-person agency use that's affordable and easy to set up." Prompts state the buyer's intent explicitly, where keywords only hint at it. Winning means getting cited in the AI's synthesized answer, not just ranking on a page.

What is Answer Engine Optimization (AEO)?

Answer Engine Optimization is the practice of engineering your content and entity signals so AI assistants like ChatGPT, Claude, Gemini, and Perplexity cite you as the answer to the questions your buyers ask. It differs from traditional SEO because the goal is inclusion in a single synthesized answer rather than a ranking among ten links. It combines clean, extractable content with trusted entity signals.

How do I find the prompts my customers ask AI about my industry?

Start with your sales calls and discovery notes, because the questions prospects ask are almost exactly the prompts they later type into AI. Mine support tickets, chat logs, site search, and community forums for real buyer language, and ask the assistants directly what people want to know when choosing your category. Then expand each seed into a full prompt with the buyer's real context and constraints.

What is the prompt space?

The prompt space is the full set of questions buyers ask AI assistants about your category, organized by buyer intent and prompt type. Mapping it turns a random pile of questions into a strategy by showing which prompts sit

closest to a purchase and which you can credibly win. You then build content and entity signals against the highest-value, most winnable prompts first.

How do I get an AI assistant to recommend my company?

You need two things: content that answers the buyer's prompt directly and cleanly, and entity signals that make the model trust you. Answer up front, structure content for easy extraction, and cover the whole prompt space, while building a consistent machine-readable identity and earning third-party mentions and reviews. Content makes you eligible to be cited; entity signals make you trusted enough to be chosen.

What is share-of-model?

Share-of-model is how often, and how favorably, AI assistants mention and recommend your brand across the prompts your buyers ask. It is the AEO equivalent of share of voice, and it replaces keyword rankings as the metric that matters. You measure it by running your priority prompts through the major assistants on a regular cadence and tracking whether you're mentioned, whether you're recommended, and what the models say about you.

How is AEO different from traditional SEO?

Traditional SEO optimizes to rank among ten links on a results page, while AEO optimizes to be cited in a single AI-generated answer that may name only a few options. SEO rewards keywords and pages; AEO rewards clean, extractable answers plus entity signals that make models trust you. They overlap, but AEO's competition for inclusion is far tighter because the answer names three options, not ten.

How often should I monitor how AI talks about my brand?

Run your priority prompts through the major assistants at least monthly, and weekly for your highest-value prompts, because answers shift as models update and the web around you changes. Record whether you're mentioned, whether you're recommended, and exactly what the model says and cites. Treat it like a dashboard you watch continuously, not a one-time audit.

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.