Scott Tischler

Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

STRATEGY

Prompt-Market Fit: How to Find the AI Questions Actually Worth Winning

By Scott Tischler · August 27, 2026 · 8 min read

The trap of the impressive prompt

When a team first discovers that AI assistants are recommending businesses, the reaction is almost always the same, and almost always wrong. They ask the biggest, most impressive question they can think of — "best marketing agency in America," "top CRM software" — see that they're not named, and set out to win it. Six months later they've spent real money moving a prompt that, even if they won it, would not have changed their business, because no actual buyer types it at the moment they're ready to spend.

I've watched this pattern enough times to give it a name. The problem isn't effort or talent. It's target selection. Answer Engine Optimization lives or dies on *which* questions you choose to fight for, and most teams choose on instinct — the prompts that sound important — rather than on evidence. The discipline that fixes this is what I call prompt-market fit.

Borrowing the most important idea in startups

Every founder knows product-market fit: the moment your product meets a market that genuinely wants it, and growth stops feeling like pushing a boulder. The lesson underneath it is that effort aimed at the wrong target is nearly worthless no matter how good the execution. A brilliant product for a market that doesn't exist fails. A brilliant answer to a question no buyer asks fails the same way.

Prompt-market fit is that idea applied to AI visibility. A prompt has fit for your business when three things are true at once:

- **Real intent** — an actual prospect types this question, in roughly these words, as part of a real decision.
- **Commercial value** — if the AI's answer named you, it would move that person measurably closer to becoming a customer.

- **Winnability** — you have a realistic path to being named, given who currently owns the answer.

Miss any one and the prompt is a bad target. High intent but no chance of winning is a waste. Winnable and valuable but nobody asks it is a trophy on a shelf. The work is finding the overlap — and being ruthless about ignoring everything outside it.

The three axes, one at a time

Intent: does a real buyer actually ask this?

This is where the keyword-era instinct betrays you. Search keywords are compressed, robotic fragments — "crm software best." Prompts are full, natural questions, often with context attached — "we're a 30-person B2B services firm outgrowing spreadsheets, what CRM should we actually use?" The second one reveals intent the first hides. If you optimize for invented keyword phrases instead of the way humans actually talk to a model, you will win prompts nobody types.

The fix is to stop guessing and start harvesting. The exact questions your buyers ask are already sitting in your sales-call recordings, your support tickets, your chat logs, and your front-line team's memory. Read fifty transcripts and you'll have a list of real prompts more valuable than anything a tool will hand you.

Commercial value: would winning it change anything?

Not all intent is equal. "What is answer engine optimization" is a real, high-volume question — and a low-commercial-value one, because the person asking is usually learning, not buying. "Who can audit my business's AI visibility" is asked by far fewer people, but every one of them is a live prospect. Rank your prompts by how close the asker is to a decision, and weight the ready-to-buy questions far above the merely curious ones. A smaller prompt with a buyer behind it beats a huge prompt with a student behind it.

Winnability: who owns the answer now?

Some answers are already locked. If a prompt is dominated by a handful of entrenched incumbents and a wall of major publisher content, a small brand's odds this quarter are low, and the honest move is to route around it. Winnability is highest where the model is currently hedging, giving generic answers, or naming brands that are weaker than you on the ground — those are the openings. Test each candidate prompt across the engines and read the answer with a cold eye: is there room for me here, or is this fortress already built?

Turning it into a scored list

Put your harvested prompts in a simple grid and score each one from 1 to 5 on the three axes: intent, value, winnability. Multiply or simply sort — the exact math matters less than the discipline of forcing the judgment. What emerges is a ranked queue. The top of that queue is your real AEO roadmap: high-intent, high-value, winnable questions where a citation would actually produce pipeline. The bottom — the impressive, crowded, low-intent prompts that started the whole exercise — is where you were about to waste six months.

Then work the queue in order. Win the fit prompts first. Each win does double duty: it produces real customers, and it builds the entity authority that makes the next, harder prompt more winnable. Authority compounds; you want to be spending it on the crowded prompts later, from strength, not throwing it at them cold today.

Why this is the leverage point

Everything else in AEO — retrievability, extractable content, off-site consensus — is execution. Prompt-market fit is target selection, and target selection is upstream of all of it. The best content in the world, aimed at a prompt no buyer types, returns nothing. Mediocre content aimed at exactly the question your ready-to-buy customer asks the model can change your quarter. I would rather a client win five prompts with genuine fit than fifty vanity prompts, and it isn't close.

One caution, offered as my read rather than a rule: fit is not permanent. The prompt that's winnable and uncrowded today attracts competitors precisely because it works, and buyer language drifts as a category matures. Re-harvest and re-score your prompt set on a cadence, the same way you'd re-check product-market fit as a market evolves. The teams that treat their prompt set as a living document, not a one-time spreadsheet, are the ones that stay in the answer as the ground shifts underneath it.

What to do this week

Pull your last fifty sales conversations or support tickets and write down the ten questions that come up most, in the customer's own words. Score each on intent, value, and winnability. Pick the single highest-scoring prompt you're not currently winning, and put this week's content and authority effort entirely behind it. That one disciplined choice — the right target instead of the impressive one — is the difference between AEO that generates a report and AEO that generates customers.

Key takeaways

- Most AEO effort is wasted chasing impressive-sounding prompts that no real buyer types or that you have no chance of winning.

- Prompt-market fit is the overlap of three things: real buyer intent, commercial value to you, and a realistic chance of being named.
- Score every candidate prompt on those three axes and work the high-intent, winnable ones first — ignore the vanity prompts.
- Harvest real prompts from sales calls, support tickets, and the questions customers actually ask, not from a keyword tool.
- A won prompt that no buyer types is a trophy; a won prompt that a ready-to-buy customer types is a pipeline.
- Re-test your prompt set on a cadence — intent and competition shift, and yesterday's winnable prompt can become a crowded one.

Frequently asked questions

What is prompt-market fit?

Prompt-market fit is the match between the AI questions you choose to optimize for and the questions that actually matter to your business. A prompt has fit when a real buyer types it, an answer that names you would move them toward a purchase, and you have a realistic chance of being named. Chasing prompts that lack any of those three is where most AEO budgets are quietly wasted.

How do I find the right prompts to optimize for?

Harvest them from reality, not from a keyword tool. Read your last fifty sales-call transcripts and support tickets, ask your front-line team what customers actually ask, and note the exact phrasing. Then test those real questions across the answer engines to see where you stand. Real buyer language beats invented search phrases almost every time.

Should I try to win every relevant prompt?

No. Some prompts are dominated by entrenched competitors or by publishers you can't outrank, and the effort is better spent elsewhere. Prioritize prompts that are high-intent and winnable, win those, and only then push into the harder, more crowded questions from a position of established authority.

About the author. Scott Tischler is the Founder & Chairman of AIrecommend.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 and professional study at Wharton, Harvard, and Oxford, he helps businesses become the ones AI recommends.