Scott Tischler

Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

STRATEGY

Topical Authority: Why Going Deep on One Subject Beats Being Everywhere

By Scott Tischler · October 14, 2026 · 9 min read

If you want AI to cite you, pick a lane and own it. Answer engines don't reward the site that has published something about everything; they reward the source they've come to read as the authority on the specific thing being asked about. Covering one subject deeply — the questions around it, the edge cases, the objections, the adjacent decisions — is what makes a model confident enough to name you. Breadth makes you findable. Depth makes you citable. Those are not the same, and most sites optimize for the wrong one.

I run AIrecommend.ai, and the pattern I see most often with businesses that aren't getting named by AI isn't that they publish too little. It's that they publish too widely. A page on this, a page on that, a blog that wanders across whatever seemed topical that month. To a human skimming, it looks productive. To a model deciding who the authority is on a narrow question, it looks like a source that knows a little about a lot — which is exactly the profile that doesn't get cited.

What topical authority means to an answer engine

Let me be precise about what I mean, because "authority" gets thrown around loosely and gets confused with author credentials. I'm not talking about who wrote the page. I'm talking about whether your body of content, taken together, reads as a comprehensive, coherent treatment of a subject.

An answer engine doesn't evaluate a single page in isolation when it's deciding who to trust on a topic. It's forming an impression of what a source is *about* and how deeply it goes there. When your content clusters tightly around one subject — the core question, the sub-questions underneath it, the common objections, the mistakes people make, the decisions that sit next to it — the model develops a confident sense that you're a source about *that*. When the model then gets a question inside that cluster, you're a natural source to pull from, because everything it knows about you points the same direction.

Contrast that with a site that has one strong page on the topic surrounded by fifty pages about unrelated things. That one page might be excellent. But it's an island. The model has no surrounding context that says "this source specializes here," so the page has to win on its own, every time, against sources that come wrapped in a whole body of corroborating depth. It usually loses.

This is the part people underestimate: depth is itself a signal. A thorough treatment of a subject tells the model you've actually grappled with it, not just touched it for traffic.

Why depth beats breadth for citations

There are a few reasons this holds, and they compound.

First, questions come in families. Nobody asks one question about a topic and stops. They ask the main question, then a follow-up, then the objection, then "okay but what about my situation." A source that has only covered the headline question can answer the first turn and then goes silent. A source that has gone deep — that has content for the follow-ups and the edge cases — can be cited across the whole conversation. In a world where a lot of the value is in the follow-up turns, the deep source gets named repeatedly while the shallow one gets named once, if at all.

Second, depth produces the specific, quotable claims that AI extracts. Broad overview content tends to stay general, because it's trying to cover a lot of ground quickly. Deep content gets into specifics — the actual number, the actual tradeoff, the actual "here's when this doesn't apply." Those specifics are what a model pulls into an answer, because they're the parts a user actually needed. Shallow content gets skimmed; deep content gets quoted.

Third, depth is a moat that breadth can't counterfeit. Anyone can publish a 600-word overview of a topic this afternoon; the internet is drowning in them, and models treat interchangeable overview content as interchangeable — which means none of it gets cited in particular. Genuinely covering the full surface of a subject, including the unglamorous edges, takes real work and real knowledge. That's exactly why it's worth something to a model trying to tell apart who actually knows the subject from who just wrote about it.

My honest read is that this is one of the most durable AEO strategies available right now, precisely because it's hard. Tactics that are easy to copy stop working the moment everyone copies them. Being the deepest source on a subject is not easy to copy, so the advantage lasts.

How to build topical authority on purpose

This is concrete enough to turn into a plan. Here's how I'd approach it.

Pick a subject narrow enough to actually own. The instinct is to pick something big because big feels like more opportunity. It's backwards. "Marketing" is not a subject you can own; "email deliverability for ecommerce" is. The narrower the subject, the fewer sources are going deep on it, and the more realistic it is for you to become the source a model reaches for. You can always widen later, after you own the first lane.

Map the question family. Before writing anything, write down every real question someone has about that subject — the main one, the sub-questions, the objections, the "how do I know if this applies to me," the mistakes, the decisions that come right before and after. This map is your content plan. It's also, not coincidentally, close to how prompts actually get phrased, which matters more than keyword lists in an AI-search world.

Cover the family comprehensively, and connect the pieces. Each question becomes content. But the pieces have to read as one body of work, not scattered posts, so link them to each other and keep the framing consistent. The goal is that a model moving through your content keeps finding the same subject treated from every angle, with the same clear point of view running through it.

Go where overview content won't. The edges are where you win: the exceptions, the "this is the part everyone gets wrong," the situation-specific advice, the honest "here's when our approach is the wrong choice." Shallow competitors skip these because they're not high-volume keywords. They're exactly what makes a deep answer useful, and exactly what a model extracts when a user's question gets specific.

Keep it current. Depth decays. A subject you covered thoroughly two years ago may have moved, and a model checking freshness will notice stale specifics. Treat your core cluster as something you maintain, not something you finish.

The mistake most sites make

The most common failure I see isn't laziness — it's diffusion. A business, wanting to grow, publishes across every topic it could conceivably be relevant to, reasoning that more surface area means more chances to be found. In a search-link world that logic had some merit. In an AI-citation world it actively works against you, because it dilutes the signal of what you're an authority on. You end up as a plausible source for everything and the definitive source for nothing, and "nothing" is what gets cited when the model is looking for an authority.

The fix is discipline, which is harder than effort. It means saying no to topics that would broaden you, and pouring that energy into going deeper on the one that's yours. It means fewer pages, not more — but each one earning its place in a coherent whole. I won't pretend that's a comfortable reallocation when every instinct says to cover more ground. But the sources AI actually names are the ones that

chose depth, and the gap between them and the sprawling, shallow sites is only going to widen as models get better at telling the two apart.

Pick the one thing you can be the best explanation of on the entire internet. Then go further on it than anyone else is willing to. That's the whole strategy, and it's durable precisely because almost no one does it.

How to tell whether you have topical authority yet

Depth is easy to claim and hard to verify about yourself, so here's a diagnostic that cuts through the self-flattery. Write out the ten most specific real questions someone would ask inside your chosen subject — not the headline question, the specific ones, including the awkward objections and the "does this apply to my situation" cases. Then check how many of them your existing content genuinely answers well.

If you can only answer two or three, you have a page on the topic, not authority over it. That gap is your roadmap — the unanswered questions are exactly where to put your next effort, because they're the follow-up turns where a deep source gets cited and a shallow one goes quiet.

The second test is simpler and more honest. Ask the major AI engines a few questions inside your subject and watch who gets named. If it's consistently the same two or three sources and you're not among them, study what they've actually published — not their homepage, their body of work. My experience is that the sources AI leans on in a category have almost always done the unglamorous thing of covering it exhaustively, and the sources that get skipped have a handful of surface-level pages and a lot of hope. The difference is visible once you look for it, and it tells you precisely what the work is: not more topics, but more depth on the one that's yours.

Key takeaways

- Answer engines cite the source they read as the authority on a specific subject, not the one that has published across the most subjects — breadth makes you findable, depth makes you citable.
- A model forms an impression of what your whole body of content is about; tightly clustered, deep coverage of one subject gives it confidence to name you, while a single strong page surrounded by unrelated pages is an island that has to win alone every time.
- Questions come in families — the deep source gets cited across the main question, the follow-ups, and the edge cases, while the shallow source gets named once if at all.
- Depth produces the specific, quotable claims AI actually extracts; broad overview content stays general

and gets skimmed rather than cited.

- Depth is a moat because it's hard to copy — interchangeable overview content gets treated as interchangeable and cited for nothing in particular.
- The common failure is diffusion: covering every possible topic dilutes the signal of what you're an authority on, leaving you a plausible source for everything and the definitive source for nothing.

Frequently asked questions

Is topical authority the same as author authority or E-E-A-T?

No. Author authority is about who wrote something — their credentials and track record. Topical authority is about whether your body of content, taken together, comprehensively covers a subject. Both help, but they're different levers: one is about the person, the other is about the depth and coherence of what you've published on a topic.

How narrow should the subject be?

Narrow enough that you can realistically become the deepest source on it, and that few others are already going deep. "Marketing" is too broad to own; something like "email deliverability for ecommerce" is ownable. The test is whether you can credibly cover the entire question family around it. You can widen later once you own the first lane.

Does this mean I should delete pages on other topics?

Not necessarily delete, but stop adding to the sprawl and concentrate new effort on your core cluster. The goal isn't fewer pages for its own sake — it's that the pages about your chosen subject form a coherent, comprehensive whole, so a model clearly reads you as the authority there rather than a generalist who touched the topic once.

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.

© 2026 Scott Tischler · scotttischler.com
scott@airecommend.ai · AIrecommend.ai