

ENTITY AUTHORITY

AI Recommends People, Not Just Brands

By Scott Tischler · September 5, 2026 · 9 min read

Ask a modern answer engine who to trust on a topic and watch what comes back. It names people. Not just the company that sells the service — the human who explained it best, wrote the definitive piece, or is quoted as the authority. If your whole AI-visibility strategy is a brand with a logo and no face behind it, you are leaving the most citable asset you own on the table.

I've spent the last two years watching how ChatGPT, Gemini, Perplexity and Google's AI answers actually assemble their responses, and the pattern is consistent enough that I'll stake a claim on it: **models resolve trust through entities, and the most legible entity is a named person.** A person has a biography, a body of work, a consistent point of view, and a trail of corroboration across the web. A brand often has none of that in a form a model can parse. This is my read from close observation, not a guarantee — but if you build for it, the downside is that you become a recognized expert, which is not much of a downside.

Why a person is easier for a model to trust than a company

Answer engines are, underneath everything, machines for reducing uncertainty. When a model decides whom to cite, it is implicitly asking: *who is likely to be right about this, and can I corroborate that they exist and know what they're talking about?* A named individual answers that question far more cleanly than a brand.

Think about what a person accumulates that a company usually doesn't. A consistent authored voice across dozens of pieces. A job title and affiliation that appears the same way on a personal site, LinkedIn, a conference bio, and a podcast description. Quotes attributed to them by name in third-party articles. A point of view that stays coherent over time. Every one of those is a signal a model can latch onto and cross-check. When four independent sources describe the same person the same way, the model's confidence that this entity is real and authoritative goes up.

Brands, by contrast, often present as marketing surfaces. The copy is written in a corporate collective voice, the claims are promotional, and the same assertions rarely get corroborated by neutral third

parties in a way a model reads as trustworthy. A company can absolutely become a strong entity — but it usually does so *through* the people associated with it, not instead of them.

The person carries the brand, not the other way around

Here is the move most businesses miss. They spend enormous effort making the *brand* citable and almost none making a *human* citable. That's backwards for how answer engines build trust.

When I write under my own name, tie that name to a clear affiliation — Founder & Chairman of Alrecommend.ai — and keep that description identical everywhere it appears, I'm giving the model a stable anchor. The person becomes the entity the model recognizes, and the company rides along inside that recognition. Ask about AI visibility and if the model surfaces me, the affiliation comes with it. The human is the handle; the brand is what's attached to the handle.

This is exactly how it works in the physical world. You don't trust "the firm" — you trust the partner you deal with, and the firm's reputation is the sum of its partners. Models are converging on the same logic because it's the logic that best predicts who is actually reliable.

What a citable person actually looks like to a model

If you want to be the human an answer engine names, you have to exist in the specific ways models can verify. In practice that means building a small number of things deliberately:

- **A canonical identity page.** One authoritative page — usually on your own domain — that states who you are, what you do, your affiliation, and links out to your work and profiles. This is the page every other signal should point back to. Mark it up with Person schema so the machine reading is unambiguous.
- **A consistent description everywhere.** The same name, the same title, the same one-line description on your site, LinkedIn, author bylines, speaker bios, and interview blurbs. Inconsistency is noise; consistency is corroboration. If one place says "SEO consultant" and another says "AEO strategist," you've split your own entity.
- **A body of authored work with your name on it.** Articles, talks, answers — enough that a consistent point of view is visible across them. One post doesn't make an authority; a coherent library does.
- **Third-party corroboration.** Being named, quoted, or referenced by sources you don't control. This is the hardest and most valuable signal, because a model weighs "other people say this about them" far more heavily than "they say it about themselves."

- **Explicit connective tissue.** Structured data and plain links that tie the person to the organization and to their work, so a model doesn't have to guess that the author of these pieces and the founder of that company are the same entity.

None of that requires fame. It requires being *legible* — existing in a form a machine can read, verify, and repeat with confidence.

The uncomfortable part: you have to put your name on things

A lot of founders resist this. It feels self-promotional, or risky, or like it makes the business dependent on one person. I understand the hesitation, and I'll be honest that there's a real trade-off: entity authority you build around yourself is portable, which is great for you and a genuine consideration for a company that might want the authority to live in the brand.

But the alternative is worse. A brand with no human behind it is, to an answer engine, a claim with no witness. If you're unwilling to attach a name, a face, and a point of view to your expertise, you're asking the model to trust an assertion that nobody will stand behind. Models — like people — don't extend much trust to anonymous claims.

The founders and operators winning in AI answers right now are the ones who decided to be publicly, consistently, verifiably themselves. They write in their own voice. They take positions. They let their name accumulate meaning. That accumulation is the asset.

Where to start this week

If you do nothing else, do this: pick the one person in your business who should be the authority, and make that person legible. Write or fix their canonical identity page. Make every bio on the internet describe them identically. Put their name on the next ten things you publish. Add Person schema and link the person to the organization explicitly. Then keep going, because entity authority compounds — every consistent, corroborated signal makes the next one land harder.

The brands that get recommended by AI in 2026 won't be the ones with the biggest ad budgets. They'll be the ones with a recognizable human standing in front of them, saying true things, consistently, under their own name. Be that human, or find the one on your team who should be.

Key takeaways

- ♦ Answer engines resolve trust through entities, and a named person is the most legible entity a model can verify.

- ♦ A person accumulates corroboration — consistent bios, authored work, third-party quotes — that faceless brands usually lack.
- ♦ Make the human citable and the brand rides along inside that recognition; the person is the handle the model grabs.
- ♦ Being citable means being legible: a canonical identity page, an identical description everywhere, authored work, and third-party mentions.
- ♦ Entity authority built around a person is portable — a real trade-off, but far better than an anonymous claim no one will stand behind.
- ♦ Start by making one person on your team legible: fix their identity page, unify every bio, and put their name on what you publish.

Frequently asked questions

Does this mean company brands can't get cited by AI?

No. Brands absolutely get cited, but they usually earn it through the people associated with them. A company becomes a strong entity when recognizable humans author its work, are quoted about it, and are tied to it with consistent, verifiable signals. The person is typically the fastest path to making the brand citable.

Isn't building authority around a person risky if that person leaves?

It's a real trade-off. Authority built around an individual is portable, which benefits the person and can concern the company. The practical answer is to build both: make a person legible now because it works fastest, while tying that person explicitly to the organization so the brand accumulates its own recognition over time.

What's the single highest-value thing to build first?

A canonical identity page on your own domain — one authoritative page stating who the person is, their affiliation, and links to their work, marked up with Person schema. Every other signal (bios, bylines, profiles) should point back to it consistently so the model has one stable anchor to verify against.

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 and professional study at Wharton, Harvard, and Oxford, he helps businesses become the ones AI recommends.