Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

ENTITY AUTHORITY

When the Web Says Two Different Things About You: How AI Resolves Conflicting Facts

By Scott Tischler · October 19, 2026 · 8 min read

When the web tells AI two different things about you — one address on your site and another on an old directory, one founding year on your About page and another in a press mention, one service list on LinkedIn and a shorter one on Google — the model doesn't flag the contradiction and ask you to clarify. It resolves it silently, and the way it resolves it is rarely in your favor. Faced with facts that don't line up, AI tends to hedge, generalize, or quietly drop the detail it can't verify. Winning here means giving every surface the same answer, so there's nothing to resolve.

I spend a lot of time looking at how answer engines describe the businesses I work with, and the single most common problem isn't that the web says nothing about them. It's that the web says *too many slightly different things*. The owner updated the website two years ago but never touched the Yelp listing. A rebrand changed the company name but three old articles still use the original one. The service page lists eight offerings; the Google Business Profile lists five; an industry directory from 2021 lists three. None of these is a lie. They're all just snapshots from different moments, and together they hand the model a jigsaw puzzle where the pieces almost fit.

What a model actually does with a contradiction

Here's the part that trips people up. We tend to imagine AI as a careful researcher who, upon finding two conflicting dates, digs until it determines the true one. That's not what happens at answer time. The model has already absorbed whatever it absorbed during training, and at the moment someone asks a question it's assembling the most probable, most defensible answer from a mix of signals — some from its training, some pulled live from a search layer. When those signals disagree, it doesn't have a referee. It has probabilities.

So it does one of a few things, and all of them cost you:

- **It hedges.** Instead of stating your founding year, it says "established in the early 2010s" or skips the detail entirely. Vague is safer than wrong, and the model is tuned to avoid being wrong.
- **It averages toward the generic.** If your service list is stated three different ways, it may describe you with the lowest common denominator — the two or three offerings every source agrees on — and omit the specialized work that actually differentiates you.
- **It picks the loudest source, not the truest one.** A contradiction between your own site and a high-traffic directory often resolves in favor of the directory, simply because that page carries more corroborating signals around it. If the directory is stale, the model quotes your old reality back to a new customer.
- **It drops you from the comparison.** This is the quiet killer. When a model is deciding who to name in a short recommendation, an entity whose facts don't cohere is a riskier citation than a competitor whose facts do. Rather than vouch for a business it can't pin down, it reaches for one it can.

Notice that none of these require anyone to have written anything negative about you. The damage comes entirely from disagreement, and disagreement is easy to create by accident.

Why your own sources are usually the culprit

When owners hear "conflicting information," they picture a competitor or a disgruntled reviewer planting bad data. In my experience that's the exception. The overwhelming majority of conflicts are self-inflicted and scattered across surfaces the business still controls but stopped maintaining.

The usual suspects: a Google Business Profile with an old suite number. A LinkedIn company page describing a version of the business from two pivots ago. A Crunchbase or industry-directory entry nobody has logged into since it was created. An old "best of" roundup that lists your previous phone number. Footers and "as seen in" badges copied between pages that quietly drift out of sync. Each one was accurate the day it was written. Collectively, they now describe three or four companies that happen to share a name.

A concrete version I see constantly: a service business moved offices eighteen months ago. The website and Google profile got updated the same week — responsible, on top of it. But the old suite number still lives in the footer of a PDF capabilities deck that ranks, in two directory listings nobody logged back into, and in a three-year-old local news write-up. Now a model answering "where are they located" has four sources, three of which quietly agree on the *wrong* answer. The business did everything right on the surfaces it thinks of as "its website," and still hands the model a majority vote for an address it left behind. Multiply that by founding year, service list and company name and you can see how a business that feels well-maintained still reads as incoherent to a machine.

The reason this matters more in AI search than it did in classic SEO is that a search engine could show ten blue links and let the human reconcile the differences. An answer engine collapses everything into one short response. It has to choose. And every unresolved contradiction is one more reason for it to choose cautiously — which, again, usually means choosing someone else or saying less about you.

How to collapse the conflict

The fix isn't glamorous and it isn't a trick. It's making every place that states a fact about you state the *same* fact, and giving the model one obvious source to anchor on. Here's the order I actually work in.

Inventory where you're described. List every surface that states facts about your business: your own pages, Google Business Profile, Bing Places, LinkedIn, Apple Maps, Crunchbase, the big directories in your category, and any press or roundups that still rank. You can't reconcile sources you haven't found. Ask the engines directly, too — query ChatGPT, Gemini and Perplexity about your name, your address, your founding year, your services — and write down every answer that's wrong or stale. That list is your repair queue.

Decide the canonical version. Pick one true value for each core fact — legal and trading name, address, phone, founding year, the real service list, the one-line description of what you do. Write them down once. This sounds trivial. It is the step most businesses skip, which is why their surfaces drift: there was never a single source of truth to copy from.

Make your own site the anchor. Put those canonical facts in plain, extractable text on your site — not locked inside an image, a logo, or a script-rendered widget — and back them with structured data (Organization and, where relevant, LocalBusiness or Person schema) so the machine-readable version matches the visible one. The goal is a page a model can read and think, *this is clearly the primary, self-published source, and it's internally consistent.*

Propagate, then prune. Update the profiles you control to match the canonical version — same name, same address formatting, same description. For stale third-party pages you don't control, request corrections where you can and, where you can't, make sure your authoritative sources so thoroughly outweigh the old ones that the stale page reads as the outlier. You rarely get everything scrubbed; you're trying to make the truth the overwhelming signal, not the only one.

Re-check, because facts rot. Consistency is not a one-time project. Addresses change, you add a service, you rebrand a product. Every change is a new chance to create a conflict, so build a quarterly pass where you re-query the engines and spot-check your top surfaces. The businesses that stay legible to AI aren't the ones that got it perfect once — they're the ones who keep the story aligned as it evolves.

What "good" looks like to a model

When you've done this well, a model researching you finds the same founding year on your site, your LinkedIn and your Google profile; the same service list, phrased consistently; the same address, formatted the same way; and structured data that confirms all of it. There's no contradiction to resolve, so the model doesn't have to hedge or hunt for a safer name. It can state who you are plainly and move on to recommending you.

That's the whole game. You're not trying to persuade the model that you're great — that's what your actual authority, reviews and track record are for. You're removing the friction that makes it *unsure*, because an unsure model is a quiet model, and a quiet model recommends the business it could describe without flinching.

My read, not a guarantee: I've watched businesses recover from being left out of AI answers simply by making their own facts agree with each other, without publishing a single new piece of content. Consistency won't manufacture authority you haven't earned. But inconsistency will hide the authority you already have, and that's a self-inflicted wound worth closing first.

Key takeaways

- When sources disagree about you, AI doesn't ask for a clarification — it silently hedges, generalizes, favors the loudest source, or drops you from the shortlist to avoid vouching for a business it can't pin down.
- Most conflicting facts are self-inflicted: a stale Google profile, a two-pivots-ago LinkedIn page, an old directory entry, a press hit with your previous phone number. None are lies; together they describe several companies that share your name.
- AI search punishes contradiction harder than classic SEO did, because an answer engine collapses everything into one short response and has to choose — and it chooses cautiously when your facts don't cohere.
- The fix is boring and it works: inventory every surface that describes you, pick one canonical value for each core fact, make your own site the consistent, structured-data-backed anchor, then propagate that version everywhere and prune the outliers.
- Consistency isn't a one-time cleanup. Every address change, new service or rebrand can create a fresh conflict, so re-query the engines and spot-check your surfaces on a regular cadence.
- My read, not a guarantee: businesses have climbed back into AI answers just by making their own facts agree, with no new content — inconsistency hides authority you already have, so close that wound

before chasing anything fancier.

Frequently asked questions

Isn't this just the old 'NAP consistency' advice from local SEO?

It's the same instinct taken further. Name, address and phone consistency was always good hygiene, but in AI search the stakes are higher and the scope is wider. It's not only your contact details that need to agree — it's your founding year, your service list, your one-line description, your positioning. An answer engine builds a single synthesized picture of you, so any fact stated two different ways across the web becomes a crack in that picture. The principle is familiar; the surface area it now covers is much larger.

How do I know which conflicting fact AI is actually using?

Ask it, the way a customer would. Query ChatGPT, Gemini and Perplexity about your name, location, founding year and services, and read what comes back. If a detail is wrong or stale, that tells you the model latched onto an outdated source — and often which one, if the phrasing matches a specific old listing. Do this before you start fixing things so you know what's actually broken, and again afterward so you can see whether the correction took. The engines are the scoreboard, not your own site.

I don't control the stale third-party page that's wrong. What then?

First, try the obvious: most major directories and profiles have a claim-or-correct flow, and it's worth the twenty minutes. For pages you genuinely can't edit — an old article, a defunct directory — you usually can't force a scrub, so the play shifts to outweighing it. Make your own authoritative surfaces so consistent and well-structured that the correct version is the overwhelming signal and the stale page reads as the obvious outlier. You're rarely going to erase every old mention; you're making the truth loud enough that the model trusts it over the noise.

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