Scott Tischler

FOUNDER & CHAIRMAN · AIRECOMMEND.AI

PLAYBOOK

Case Studies Are Citation Assets: Writing the Proof AI Quotes Back

By Scott Tischler · October 12, 2026 · 8 min read

A case study is one of the few things you can publish that gives an AI model exactly what it needs to recommend you with confidence: a specific, named, measurable outcome tied to a specific situation. The problem is that most businesses either bury that proof inside a PDF no crawler reads, or write it in vague, self-congratulatory language a model can't lift. Written for extraction instead of for a sales deck, a case study becomes a quotable piece of evidence — the kind answer engines repeat back to the person asking about your category.

I run AIrecommend.ai, and when I audit why a business isn't getting recommended, a recurring pattern is that the company has done genuinely impressive work and has almost nothing citable to show for it. The results are real. They're just trapped — in a slide the sales team emails, in a client's private testimonial, in a number the founder quotes on calls but never wrote down anywhere a machine can see. The work happened. The evidence didn't get published in a form that counts.

This piece is specifically about case studies as an answer-engine asset. It's related to publishing original data, but it's not the same thing — a data study is proprietary research you run; a case study is the story of a real outcome you delivered. Both earn citations, but they earn them differently, and the case study is the one almost every business already has the raw material for and simply hasn't packaged.

Why a case study beats a testimonial for AI

A testimonial says "they were great to work with." A model can read that, but it's weak evidence — it's an opinion with no specifics, and the web is saturated with it. A case study says "a 40-person logistics company cut their quote turnaround from three days to four hours after switching, and closed 22% more of the quotes they sent." That sentence contains a who, a situation, a mechanism, and a measurable result. When a model is deciding whether to recommend you for "fast quoting for logistics firms," the second sentence is usable and the first one isn't.

The difference comes down to specificity, which is one of the factors that turns a hedged AI mention into a confident one. General praise supports a general impression. A specific, named outcome supports a specific query — and specific queries are exactly where smaller businesses can win, because that's where the national brands haven't bothered to accumulate proof. A case study is a specificity machine. It's your chance to tie your name to a precise job, a precise customer type, and a precise result, all in language a model can connect to a real question.

What makes a case study quotable

Not every case study earns citations. The ones that do tend to share a few traits, and they're almost all about being concrete enough to extract.

A real number, stated plainly. "Significantly improved efficiency" is noise. "Reduced no-show rate from 18% to 6% over four months" is signal. You don't need a spectacular number — you need a specific one. A model can repeat "cut costs by 15%" far more usefully than "delivered major savings," and a modest honest number beats an impressive vague one every time.

A named, recognizable situation. You don't have to name the client if they won't allow it, but you do have to name the situation precisely: the industry, the size, the specific problem. "A regional HVAC company with eight trucks" is extractable context. "A valued client" is not. The more clearly the situation maps to the kind of customer you want AI to send you, the more useful the proof is.

The mechanism, not just the result. The model — and the human reading the model's answer — wants to know *how*, because that's what signals the result is repeatable rather than luck. One or two sentences on what you actually did turns a result into a method, and a method is what gets recommended.

It's on a crawlable web page. This is the one that quietly kills most case studies. If your proof lives in a gated PDF, a slide deck, or a "request the full story" form, no answer engine will ever read it. The case study has to be an actual HTML page, publicly accessible, with the key facts in the body text — not locked away and not trapped inside an image.

The structure a model can actually lift

You can write a case study a hundred ways, but the ones that extract cleanly tend to front-load the facts in a predictable shape. I'd lead with a one-sentence summary that contains the whole story: who, what changed, and by how much. Put that sentence near the top in plain prose, because that's the sentence a model is most likely to quote directly.

Then give the situation before, the action you took, and the result after, each in its own short, clearly labeled section. The "situation before" establishes the problem and the customer type. The "what we did" establishes the mechanism. The "result" delivers the number. A model reading top to bottom can pull any of those three depending on what the person asked, and a human skimming gets the story fast too. You're not choosing between writing for machines and writing for people here — the structure that extracts cleanly is also the structure that reads well.

Resist the urge to wrap it all in marketing language. "We're thrilled to have partnered with an amazing client on an exciting journey" is pure filler, and a model trying to find the actual facts has to wade through it. Say what happened. The restraint reads as credibility, and credibility is the whole point of proof.

The mistakes that keep good results uncitable

Most businesses don't have a proof problem — they have a packaging problem. The results exist; they're just published in ways that guarantee no model will ever use them. A few patterns come up over and over, and they're all fixable in an afternoon.

The first is hiding the number behind a verb. "Dramatically increased bookings" feels strong to write, but it's unusable — there's nothing to extract. Go find the actual figure and state it. The second is burying the whole thing in a carousel of logos or a testimonial slider, where the words exist only inside images and a crawler sees nothing. If the facts aren't in real text, they don't count. The third is the gate: locking your best proof behind a form because someone decided case studies were lead magnets. A gated case study generates a handful of email addresses and zero citations; an open one works for you around the clock. The fourth is writing the study as a press release about how thrilled everyone is, so the one real sentence of substance is drowned in congratulation. Trim the celebration and let the facts stand — a model, like a skeptical buyer, is looking for the result, not the adjectives.

Where case studies fit in the bigger evidence picture

A single case study on your own site is still your claim about yourself — valuable, but one witness. The real leverage comes when the outcome gets corroborated elsewhere: the client mentions it in their own content, a trade publication writes it up, you reference it in a podcast or a guest article, a review describes the same result in the customer's own words. Now the same specific outcome is attested from several independent directions, and that independence is what earns the confident recommendation instead of the hedge.

So I think about case studies as seeds rather than endpoints. You publish the clean, citable version on your site, and then you work to get that same story echoed by people who don't work for you. One

strong outcome, corroborated across three or four independent places, does more for your AI visibility than a dozen vague testimonials sitting on a wall. The case study gives the proof a canonical home; the corroboration gives it weight.

How to publish one this week

You almost certainly already have the material. Pick your best recent result — the one where a specific customer got a specific, measurable outcome you'd be proud to have quoted. Write the one-sentence summary first and make yourself put a real number in it. If you don't have the number, this is the week to go get it from the client; that number is the most valuable sentence you'll publish this quarter.

Then draft the situation, the action, and the result as three short sections in plain language, publish it as a normal public web page — not a PDF, not a gated download — and link to it from your site so it gets discovered and crawled. Finally, do the unglamorous follow-up: ask the client if they'll reference it, pitch the story to one relevant publication, and work it into the next podcast or article you contribute to. That last step is what turns a page on your site into evidence across the web.

My honest read, and it's a read rather than a promise: you won't see a recommendation flip the week you publish. This compounds. But case studies are among the highest-return things you can publish for AI visibility, because they're the rare asset where you control the proof, the specificity is natural, and the raw material is work you've already done. Most of your competitors have the same results sitting in a drawer. The ones who publish them in a form a machine can read are the ones who get quoted.

Key takeaways

- A case study hands an AI model what a testimonial can't: a specific, named, measurable outcome tied to a specific situation — the exact kind of proof that supports a specific recommendation.
- Quotable case studies share four traits: a real number stated plainly, a named situation, the mechanism behind the result, and a crawlable public web page — not a gated PDF or an image.
- Lead with a one-sentence summary containing who, what changed, and by how much; it's the sentence a model is most likely to quote directly.
- A case study on your own site is one witness; its power multiplies when the same outcome is corroborated by the client, a publication, a review, or a podcast mention.
- Marketing filler actively hurts extraction — say what happened plainly, because the restraint reads as credibility and lets a model find the facts.

- You already have the raw material; the edge isn't doing the work, it's publishing the proof in a form a machine can read and then getting it echoed independently.

Frequently asked questions

Do I have to name the client for a case study to help my AI visibility?

No. You can keep the client anonymous as long as you name the situation precisely — the industry, the size, the specific problem, and a real result. 'A regional HVAC company with eight trucks cut its no-show rate from 18% to 6%' is fully extractable without naming anyone. What a model needs is the specific situation and the specific number, not the brand name.

Is a case study the same as publishing original data for AEO?

They're related but different. Original data is proprietary research you run — a survey, a benchmark, an analysis — and it earns citations as a source of facts. A case study is the story of one real outcome you delivered, and it earns citations as proof of capability. Most businesses already have case-study material from work they've done, whereas original data requires running the research. Both are worth doing; the case study is usually the faster win.

Why won't my case study get cited if it's in a PDF?

Most answer engines either don't read gated or linked PDFs reliably, or weight them far below normal web pages, and a PDF behind a form is invisible to them entirely. If your proof lives in a downloadable deck or a 'request the full story' gate, treat it as unpublished for AI purposes. Put the key facts in the body text of a public HTML page, and keep the PDF as an optional extra for humans who want it.

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