

EBOOK · 2026

How scaling B2B brands win in *AI Search*.

How scaling B2B brands win in search and earn citations and preference from ChatGPT, Perplexity, and Google's AI Overviews.

Who this is for. Marketing leaders at scaling B2B brands who want to understand what is happening in search, what the evidence says about AI visibility, and how Eleven approaches the work.

01 At a glance

Everything in this guide is the evidence behind five statements.

- 1 Your buyers now research vendors inside AI tools.** Forrester's 2026 survey of nearly 18,000 business buyers found 94% used generative AI during their most recent purchase. It now ranks as their most meaningful research source, ahead of vendor websites, product experts, and sales reps, and the number of buyers saying so doubled in a year.
- 2 The shortlist forms before you know the deal exists.** G2's March 2026 research found 51% of B2B software buyers now start their research in an AI chatbot such as ChatGPT more often than in Google, up from 29% a year earlier.
- 3 SEO is fundamental.** Across the brands we track, AI visibility and search performance move together: the pages models cite are very often the ones already earning impressions and clicks in Google, and the same work that drives SEO performance also drives AI visibility.
- 4 Most AI citations come from earned media.** Muck Rack's analysis of more than 1 million AI prompts found approximately 84% of unpaid citations originate from earned coverage rather than brand-owned pages. Coverage starts with expert content worth covering, which is why Digital PR is the add-on we recommend alongside every content engagement.
- 5 The work compounds.** Expert content, brand mentions, and presence across the channels AI systems read reinforce each other. Eleven delivers all three, or whichever parts a team needs.

02 How buyers research vendors now

Turn G2's and Gartner's buyer surveys into a single person and the change is easy to see. A Head of Operations at a 150-person company needs time-tracking software. Two years ago they opened Google, clicked through comparison articles, and built a spreadsheet.

Today, they are more likely to start by asking ChatGPT for three options suited to a company her size, then push back with follow-up questions until the list feels right. She will still read the reviews, visit the vendor sites, and run the demos and committee scrutiny that decide the deal. What has changed is where brands first enter her awareness.

A process that began with ten open tabs now begins with one question and one answer, and the brands in that first answer are the ones she researches from that point on.

69%

of B2B software buyers chose a different vendor than planned based on AI chatbot guidance, and one in three bought from a vendor they had never heard of before the chatbot named it (G2, The Answer Economy, March 2026).

Gartner's 2026 study of 645 buyers cuts the other way: seven in ten still validate AI-generated insights with a sales rep before committing. The influence of AI search is concentrated at the start of the journey, where it shapes which brands are considered, while the later stages of the deal still rely on human judgement.

What this means for a scaling B2B brand. The shortlist was always drafted before a vendor ever heard from the buyer; what changed is that a model now drafts it, from sources you influence rather than pages you control. If the model cannot see you, neither can the buyer.

03 The new shape of search

Search is not disappearing; it is reshaping, and the impact on B2B SaaS marketing is already measurable in click-through rates, referral behaviour, and the formation of buyers' first impressions of a brand.

Impressions are up, clicks are down. BrightEdge tracked a roughly 49% rise in impressions since AI Overviews launched, alongside a 30% drop in click-throughs. Ahrefs reported a 34.5% fall in referrals to ranking pages. Visibility now happens upstream, inside the AI summary itself.

The traffic you do get is more valuable. Buyers who reach your site after an AI conversation have already compared their options inside it, and they behave differently once they arrive. Similarweb's June 2026 panel study tracked users who received a ChatGPT brand recommendation and found they viewed 12.0 pages and spent 11.8 minutes on site, against 6.5 pages and 5.6 minutes for other visitors.

Exposure Ninja's March 2026 analysis found AI search traffic converting at 14.2%, compared with 2.8% for standard organic visits, and Eleven's partner data shows the same pattern for demos and sign-ups.

The lost clicks still matter: a randomised field experiment found that the visits AI Overviews suppressed on Google were no lower in quality than those of other visits (Agarwal and Sen, 2026).

Brand discovery is moving upstream. A Siege Media study of 50 websites recorded total traffic down 5% while homepage traffic rose 10.7%. Buyers encounter brands in AI-generated answers and then go directly to websites. Branded search is now a leading indicator of demand.

PLAIN ENGLISH: THE TERMS IN THIS GUIDE

Impression. One appearance of your page in a set of search results, whether or not anyone clicks it.

Referral. A visit that arrives from another platform, such as a link inside a ChatGPT answer.

Branded search. A search that includes your company or product name, a strong signal that someone already knows you exist.

Citation. A link or named source inside an AI answer; how AI visibility becomes traffic and pipeline.

04 GEO, AEO, AIO: it's SEO + Brand Marketing

The industry has spent two years debating names for optimising content in AI search, and the acronyms all describe roughly the same work. The evidence points one way: the strongest predictors of AI citations are the same signals that drive search performance.

In May 2026, Google made its position official, saying that optimising for AI Overviews and AI Mode is "still SEO". Google can only speak to its own features, which is why we also ran our own correlation study across ChatGPT, Perplexity, and Claude, and found that visibility in LLMs and Google Search Console impressions correlate closely.

This has been Eleven's position from the start, and it is why we treat SEO and AI Visibility as one practice rather than selling a new acronym. Google's guidance also names the content it most wants for the AI era: material with a unique point of view, grounded in first-hand experience. That is a description of expert-led writing.

One weighting does change, though. Brand popularity, authority, and sentiment carry more influence on AI visibility than on classic search rankings: in Ahrefs' study of 75,000 brands, branded mentions were the strongest correlate of AI Overview visibility. Models build their picture of a brand from what the wider web says about it, which is why the brand-building recommendations later in this guide matter as much as the on-page ones.

What you can safely deprioritise

Google is equally direct about tactics that do not work for its AI features, several of which are being sold hard right now.

- ✗ **Publishing lms.txt files.** Google Search does not currently use them, and while they cost little to implement, nobody should be selling them as a visibility strategy.
- ✗ **Chopping content into machine-sized fragments.** Google advises writing for people and letting its systems handle the extraction.
- ✗ **Rewriting existing pages specifically for AI engines.** There is no special dialect that AI systems prefer, so clear, well-sourced writing wins on both surfaces.
- ✗ **Buying or manufacturing brand mentions.** Google treats inauthentic mention-building the way it treats link schemes, leaving earned coverage as the only durable form.

Some structural good practices still help retrieval in other LLMs, such as ChatGPT and Perplexity, like answering the question early on the page and supporting claims with named sources. Still, none of it substitutes for the authority signals this guide is about.

05 How AI search works, briefly

AI systems learn about brands through two routes that work together. Training data builds durable associations: when a brand appears consistently across high-authority sources, the model learns to connect it with its category, which is why brand work compounds.

Live retrieval, also known as retrieval-augmented generation (RAG), brings in up-to-date information: the chatbot searches the web while you wait and assembles a cited answer based on what it finds.

TRAINING DATA

Durable associations

Consistent presence across high-authority sources teaches the model to connect a brand with its category, ahead of any single conversation.

LIVE RETRIEVAL (RAG)

Up-to-date answers

The chatbot searches the web in the moment and assembles a cited answer from what it finds — shaped by the associations training already formed.

Training shapes retrieval too: the model writes its searches and weighs what comes back using the associations it has already formed, so the two routes reinforce each other rather than dividing the work.

The pages retrieval pulls in are very often the pages search engines already rank, but not only those. LLMs cast a wider net than the top ten, and a page that answers one narrow question well can be cited from far deeper in the results, which is an opening for brands not yet appearing on page one on search.

When you ask a chatbot a question, it rarely searches for your exact words. Instead, it breaks your prompt into several narrower searches, runs them in parallel, and blends the results, a process known as "query fan-out". Peec AI's analysis of 20 million of these behind-the-scenes ChatGPT queries found their average length roughly doubled between October 2025 and January 2026, from about 6 words to about 12. The machine's questions are getting longer and more specific, which rewards brands whose content covers a topic in genuine depth over brands with one thin page per keyword.

06 What actually predicts AI citations

Claims about AI visibility are easy to make and hard to check, so we ran the numbers ourselves. Eleven's AI Citation Correlation Study matched citation data from ChatGPT, Perplexity, and Claude against Google Search Console performance, link profiles, and page characteristics for approximately 300 pages and 1,257 URLs from a single B2B partner. It is a small, directional study, and we are expanding it as more partner data arrives, but the direction held across all three engines, though the strength varies by platform, and it matches the far larger public studies below.


What predicts AI citations

Correlation (r) with how often a page is cited by ChatGPT, Perplexity, and Claude

STRONG PREDICTORS



WEAK PREDICTORS



Eleven AI Citation Correlation Study: ~300 pages, 1,257 URLs, across ChatGPT, Perplexity, and Claude.

Search visibility predicts AI citations. *The pages Google shows are very often the pages LLMs cite.* The strongest predictors aggregated across the three engines were Google Search Console impressions (r=0.82), referring domains (r=0.79), total backlinks (r=0.79), GSC clicks (r=0.76), and the breadth of keywords a page ranks for (r=0.74). Citations correlated with referral traffic from AI platforms (r=0.84): the pages models cite are also the pages receiving real visits.

The weaker signals are just as instructive

Word count barely registers at r=0.18, because pages earn citations through depth, first-hand expertise, and verifiable sourcing rather than sheer volume. That is E-E-A-T doing the work, and it is why thin pages rarely appear in AI answers, whatever their word count.

Internal linking has almost no direct effect on citations ($r=0.09$), though it still matters for crawlability and rankings, so we treat it as an SEO job rather than an AI tactic. Anyone selling internal linking as an AI visibility lever is reading the data selectively.

Link diversity beats link volume: for ChatGPT, referring domains ($r=0.78$) predicted citations more strongly than total backlinks ($r=0.73$), so mentions from a wider set of sources beat a taller pile from the same few places, a finding about PR strategy as much as link building.


The platforms also behave differently from each other. In our data, ChatGPT's citations follow conventional search signals closely (impressions $r=0.81$), while Perplexity's follow them far more loosely (impressions $r=0.61$, referring domains $r=0.39$) and lean more on community and review sources. Averi's March 2026 analysis of 680 million citations found that only 11% of domains were cited by both platforms, so visibility on one platform does not automatically transfer.

Larger studies show the same pattern

AirOps' March 2026 analysis of 548,534 pages retrieved by ChatGPT found that 43.2% of pages ranking first in Google were cited, 3.5 times the rate of pages outside the top 20, and that 55.8% of all cited pages ranked in Google's top 20 for at least one query. At the brand level, Ahrefs' analysis of 75,000 brands found branded web mentions had the strongest correlation with AI Overview visibility (0.664), ahead of branded anchor text (0.527) and branded search volume (0.392). AI mostly cites the pages search already surfaces and recommends the brands the web already talks about.

Ranking in Google predicts ChatGPT citations

Share of pages cited by ChatGPT, by Google ranking position



55.8% of the pages ChatGPT cited ranked in Google's top 20 for at least one query.

AirOps, March 2026: 548,534 pages retrieved by ChatGPT across 15,000 prompts.

07 Where LLMs get their evidence

If citations drive visibility, the next question is where citations come from. As it turns out, the answer reshapes how a content budget should be spent.

- 1** Approximately **85% of unpaid AI citations originate from earned media rather than brand-owned pages** (Muck Rack's analysis of more than 25 million AI citations, May 2026). Models cite publications they already consider credible far more often than vendor websites, though the mix depends on the question: Muck Rack found that industry trend prompts cite journalism at more than double the rate of how-to and comparative prompts. For a B2B brand whose buyers ask comparison questions, earned media has to stretch beyond press coverage into the review sites and communities those prompts draw on.
- 2** **The most-cited domains across major models are YouTube, Reddit, Wikipedia, NIH, and Statista** (Meltwater's analysis of 8 million citations across eight LLMs, May 2026), and earned coverage remains among the most influential visibility drivers in that dataset.
- 3** **LinkedIn moved from outside the top 20 to the most-cited domain for professional queries** between November 2025 and February 2026 (Profound). What your leadership publishes there feeds directly into how models describe your company.

The economics behind this pattern dates back to Sherwin Rosen, who showed in 1981 that when the preferred option can reach a large market at low marginal cost, small differences in perceived preference can produce large differences in reward. LLMs push that logic further because an answer names only a handful of brands, and the brands that keep appearing across relevant, credible sources become the corroborated choice, the one the model finds again wherever it looks. Peer-reviewed work adds a caveat: retrieval systems respond to a tangle of stylistic and provenance cues rather than a single clean authority signal, which is one more reason to spread presence across many credible sources. Authority accrues from the outside, through the references, mentions, and citations the wider network supplies.

Your own site is a small part of the picture. The model's view of your brand is assembled mostly from sources you influence rather than those you control: trade press, review platforms, communities, and professional networks. Brand work and Digital PR sit inside our search offering for exactly that reason, and audio and video carry the same weight because they are the formats that travel furthest across the channels models read, with YouTube sitting at the top of Meltwater's most-cited list.

08 What AI says about you, against your competitors

Visibility is only half the picture, because when a buyer asks a chatbot to compare vendors, the answer names brands and frames each one. That framing works like a rating in the buyer's mind, even though no scoring occurs beneath the surface. The same brand can read differently from one query or model to the next, and the description is what the buyer carries into the shortlist conversation. That framing has a real effect on buying decisions, so influencing how models describe your product is one of the strongest levers on purchase intent. G2 found that 85% of buyers think more highly of a vendor a chatbot recommends, and G2 measured only the upside: the mirror image follows on its own, because a recommendation delivered with caveats carries those caveats into the shortlist conversation.

At Eleven, we assess how the main models describe a brand relative to its competitors across the factors buyers in that category care most about. For a scaling B2B brand, that typically means ease of implementation, product depth, integrations, pricing transparency, support quality, and security and compliance.

Each gap that assessment surfaces points to its own fix: a weak showing on pricing transparency is usually a content problem, where the information models need is missing, thin, or buried behind a demo request. A weak showing on support quality tends to be an evidence problem, where reviews and community discussion point the wrong way, which is reputation and PR work rather than on-page work. Reading the gaps this way tells you which lever to pull, and looking again once the work has landed shows whether the models' picture changed.

The models also disagree with each other more than most teams expect, because each reads a different diet of sources: Perplexity leans on communities and reviews, ChatGPT follows search visibility more closely, and Gemini draws on Google's index. A brand can be the recommended option in one model yet barely appear in another, mirroring the 11% citation overlap noted earlier in this guide. One strategy still serves all of them: telling the same story across more source types covers every platform at once, and examining models individually reveals which sources are missing.

09 How to win the new search

Three levers drive AI visibility, and the scaling B2B brands performing in AI search today are investing in all three. Together they are Expertise, Everywhere in practice: genuinely expert content, distributed across the owned and earned channels that buyers and models both read.

01

FOUNDATION

Expert-led content

Content that gets cited.

- Named experts in the byline
- E-E-A-T: experience, expertise, authoritativeness, trust

02

MULTIPLIER

Brand mentions and Digital PR

Mentions that compound.

- Digital PR in credible publications
- Branded mentions feed AI training data and live retrieval

03

COMPOUNDER

Presence across the channels AI reads

Presence that adds up.

- LinkedIn, YouTube, communities, audio
- One asset, many formats; repetition builds associations

1. Expert-led content. Content from named subject-matter experts, carrying the E-E-A-T signals that Google's quality systems and LLM retrieval both reward: demonstrable experience, expertise, authoritativeness, and trust. It is what Google's guidance calls for by name, and it is the one input AI-generated content cannot fake.

2. Brand mentions and Digital PR. Original research, expert commentary, and story-led outreach that earn coverage in the publications LLMs draw from. Branded mentions are the strongest brand-level predictor of AI visibility in Ahrefs data, and earned coverage is where unpaid citations originate.

3. Presence across the channels AI reads. A small number of well-chosen platforms done properly: for B2B that increasingly means LinkedIn, YouTube, relevant communities, and audio, alongside the website.

Each lever propels the other two: expert content gives PR a story worth pitching; coverage earns the mentions and branded searches that models weigh most heavily; and presence across a few channels puts every piece of work in more of the places AI systems read. Audio sits in the stack for exactly this reason: podcasts put founders and operators on the record in a format models increasingly draw from, and YouTube now sits among the most-cited domains across every major LLM.

Underneath all three, content works like a heartbeat: models reward a steady pulse of expert material across the sources they read, and a stockpile that stops updating fades from answers. AI collapsed the cost of producing generic content to almost nothing, which is exactly why generic content stopped working. The cost of genuine expertise has not fallen, and if anything, it has risen, because experts are in greater demand now that generated content is everywhere. That is why quality runs throughout our process, from the idea to the expert input to the finished piece. Our writers carry an MA or equivalent industry qualification, three or more years of real-world experience, and three or more years of bylined journalism in the sectors they cover, which is the vetting bar that makes content worth citing.

10 Where Eleven fits

Eleven was built for the way search is changing: we work with scaling B2B brands that need to be found, cited, and recommended by human buyers and machine retrieval systems, under one principle: Expertise, Everywhere. The work runs through three connected disciplines, underpinned by reporting.

Strategy	Market research, SEO and AI visibility audits, content strategy, and the technical foundations that AI citations depend on. Every recommendation connects content to business goals rather than working from a fixed checklist.
Production	Vetted topic experts matched to your sector, your own internal expertise drawn out through interviews and review cycles, and professional editors applying a six-pillar quality framework to deliver publish-ready written, graphical, and audiovisual content built to be cited and trusted.
Amplification	Digital PR and story-led outreach that earn coverage in the publications LLMs draw on, podcasts and video that put your experts on the record, newsletters, and the social and community engagement that keeps a brand visible between site visits.

Reporting and analytics underpin all three, through monthly reporting and quarterly reviews that keep the strategy tied to results rather than activity.

How we work

- ✓ **We lead with expertise.** Every piece of content comes from or is shaped by a genuine subject-matter expert.
- ✓ **We build fame.** Success is whether you are becoming more known, more trusted, and more preferred in your category, and traffic follows from that.
- ✓ **We stay grounded in evidence.** Recommendations come from research, our own data, and rigorous testing. We are honest about what we know and what we do not.
- ✓ **We focus on what compounds.** Brand associations, authority signals, thought leadership positions, and community trust pay off across years, not months.
- ✓ **We do not chase gimmicks.** There is no shortcut to authority, and anyone selling one is selling a problem you will have to clean up later.

How we measure success

Measurement in AI search is imperfect, and we would rather say that than sell a dashboard that pretends otherwise. Third-party LLM tracking is directional rather than definitive, and attribution across AI platforms remains fragmented. Within those limits, we track share of search, buyer-intent rankings and traffic, branded search volume, LLM visibility and sentiment across the main platforms, and LLM referral traffic and conversions in GA4. We do not chase vanity traffic, content volume as a KPI, or rankings with no commercial or brand-building value, and we share what our own research finds with partners as the picture sharpens.

Where to start. We can run an AI citation correlation study on your domain, audit your SEO foundation, and map the brand mentions your category requires to compete. A 20-minute strategy call is the fastest way to find out whether we are a fit.

11Content.com

11 Sources

1. [Forrester, Buyers' Journey Survey, ~18,000 buyers \(2026\)](#)
2. [G2, The Answer Economy, 1,076 buyers \(March 2026\)](#)
3. [Gartner, B2B buyer survey, 645 buyers \(2026\)](#)
4. [Agarwal and Sen, AI Overviews field experiment \(SSRN, 2026\)](#)
5. [BrightEdge, AI Overviews impact tracking \(2025\)](#)
6. [Ahrefs, click-through analysis \(2025\)](#)
7. [Ahrefs, 75,000-brand mentions study \(2025\)](#)
8. [Similarweb, The Downstream Impact of AI Visibility \(June 2026\)](#)
9. [Exposure Ninja, AI search conversion analysis \(March 2026\)](#)
10. [Siege Media, 50-website traffic study \(2025\)](#)
11. [Google Search Central, generative AI guidance \(2026\)](#)
12. [Peec AI, 20 million query fan-outs \(2026\)](#)
13. [Muck Rack, What is AI Reading? \(2026\)](#)
14. [Meltwater, GenAI Lens, 8 million citations \(May 2026\)](#)
15. [Profound, platform citation research \(2026\)](#)
16. [Averi, 680 million citations analysed \(March 2026\)](#)
17. [Rosen, The Economics of Superstars \(1981\)](#)
18. [Eleven Content, AI Citation Correlation Study \(dashboard on request\)](#)