

BRIDGEHEAD

A BRIDGEHEAD COMMUNICATIONS EVIDENCE REVIEW • SEPTEMBER 2026

RETRIEVED, RANKED, **CITED**

An evidence-graded review of 350 sources on what actually moves organic rankings and AI-answer citation in 2026, and what it means for one real website

“

RETRIEVAL IS THE BOTTLENECK, NOT PHRASING.

CONCLUSION TWO · RETRIEVED, RANKED, CITED, 2026



TABLE OF CONTENTS

WHAT IS IN THIS REPORT

-	How to read this	4
01	Ten conclusions	5
02	What changed, 2024 to 2026	7
03	What moves organic rankings	9
04	What moves AI citation	11
05	What buyers and journalists respond to	14
06	How to structure and write an Insight	16
07	Technical plumbing	19
08	Measurement	21
09	A worked example: one site today	23
10	The programme	26
11	Contested points and sources	28

A NOTE ON THIS REPORT

Published by Bridgehead Communications, 2 September 2026, and reissued in this designed edition in October 2026 with the text of the published review unchanged. The per-claim research dossiers behind each grade, with URLs, dates and sample sizes, are available on request. Where this report and a dossier differ in grading, the more conservative grade applies; in particular the Limy blog post is graded as a commercially motivated summary of the Martinez survey, not as peer-reviewed work in its own right.

INTRODUCTION

HOW TO READ THIS

The brief was to cut through unevidenced speculation about SEO and GEO. So every substantive claim below carries an evidence grade. The grade describes the quality of the evidence, not the size of the effect. Where two credible sources disagree, both are given. Where a

widely repeated figure could not be traced to a primary source, it is marked unverified rather than dropped, because knowing which "facts" are folklore is half the value of the exercise.

- [A] Primary document or peer-reviewed work. Google and Microsoft documentation, court findings of fact, refereed papers, independent panels such as Pew.
- [B] Large sample with a credible design, even if vendor-published. Quasi-experiments, server logs, clickstream panels, controlled benchmarks.
- [C] Large sample, vendor-published, correlational. Reliable for direction, not for causation or magnitude. Most Ahrefs, Semrush, Seer and BrightEdge studies sit here.
- [D] Small sample, single site, agency case study or expert opinion. Useful for hypotheses only.
- [U] Circulates widely but the primary source could not be reached or does not say what is claimed.

Method. Six parallel research streams (official Google and Bing guidance, empirical generative-engine research, organic ranking evidence, B2B thought-leadership research, content structure and writing, technical plumbing and measurement) each produced a graded dossier from primary sources. The worked-example baseline in section 9 was pulled live on 2 September 2026 from the Ahrefs API (Search Console passthrough, AI citation counts, domain metrics), Bright Data UK search results, and the site's own source code. Sources fetched are listed in section 11.

EXECUTIVE SUMMARY

TEN CONCLUSIONS

1. **There is no separate GEO discipline for Google, by Google's own account.** Its May 2026 optimisation guide, its December 2025 AI-features documentation and three named spokespeople all say the same thing. AI Overviews and AI Mode draw on the core ranking systems, need no special files, markup or chunking, and Google Search ignores llms.txt. [\[A\]](#)
2. **Retrieval is the bottleneck, not phrasing.** The peer-reviewed record (C-SEO Bench at NeurIPS 2025, the SAGEO Arena preprint, the Martinez survey of 45 studies) finds that citation-oriented rewrites are mostly null in realistic pipelines and can reduce a page's chance of being retrieved at all. The famous "30 to 40 percent uplift" from the 2024 GEO paper applies only to a page already in the model's context. [\[A\]](#)
3. **Once retrieved, two things reliably help.** Close relevance between the page (and its headings) and the literal question, and extractable evidence in the passage: a definition, a number with a source, a named quote, a date. These replicate across the academic and industry data. [\[A\]](#) [\[C\]](#)
4. **User satisfaction is the best-evidenced organic ranking input.** The 2024 antitrust findings of fact establish that Navboost memorises 13 months of click, dwell and return-to-results behaviour per query and was not replaced by language-model signals. Pages that end the search win. [\[A\]](#)
5. **Links matter less than folklore says, and brand mentions matter more.** Google removed "important" from its description of links in March 2024. In the UK results checked today, pages with zero referring domains rank in the top ten for sector-qualified agency terms. Across 75,000 brands, unlinked web mentions correlate with AI visibility three times as strongly as backlinks. [\[A\]](#) [\[C\]](#)
6. **AI Overviews have cut informational clicks by a third to a half, but branded and commercial queries are largely spared.** Pew's independent panel, Seer's 2.4 billion impressions and two Ahrefs studies agree on direction. Being cited roughly doubles the residual click rate. [\[A\]](#) [\[C\]](#)
7. **Each engine has its own source ecology and they barely overlap.** Any two engines share about 17 percent of cited domains, and roughly two thirds of an answer's sources change from one day to the next. Single-run "are we cited" checks are noise. [\[C\]](#)
8. **Earned coverage dominates non-Google citation.** Around 84 percent of ChatGPT, Claude and Gemini citations go to earned media, and third-party pages out-cite owned pages on ChatGPT and Perplexity. For a PR firm this is the home fixture. [\[C\]](#)
9. **Buyers reward original research and named experts, and punish commodity content.** Three quarters of decision-makers have researched a firm they were not considering because of its thought leadership; 86 percent want to be challenged rather than reassured; in-house technical experts are 22 points more credible than chief executives. Only 15 percent rate what they read as very good. [\[C\]](#)
10. **In the worked example in section 9, the binding constraint was not authority, it was snippet integrity and content, in that order.** The site already ranked first for its primary commercial term and converted at 0.07 percent, because the page's title and description were truncated in source and Google was substituting schema boilerplate. Fixing that is the cheapest measurable win available on a site in that position. Section 9 sets out the evidence. [\[A\]](#)

01

PART ONE · SECTIONS 2 TO 4

THE EVIDENCE

What changed between 2024 and 2026, what the primary record says moves organic rankings, and what moves citation in AI answers. Every claim carries its grade.

COVERS

The click economy, platform guidance, the trial record and the leak, the citation pipeline

KEY THEME

Retrieval and user satisfaction are the best-evidenced inputs; most GEO tactics are not

PART I • SECTION 2

WHAT CHANGED, 2024 TO 2026

The click economy

8% VS 15%

Share of Google visits ending in a click on a result when an AI summary is present vs absent. Pew, 900 US adults, 68,879 tracked searches, Jul 2025. [A]

68%

US searches ending without a click to the open web, Jan to Apr 2026 (Similarweb panel; 2024's 58.5% used a different panel). SparkToro, Jun 2026. [B]

-58%

Position-one organic CTR on informational queries with an AI Overview, Dec 2023 to Dec 2025, 300k keywords. Ahrefs, Feb 2026. [C]

1.08%

Share of enterprise web sessions referred by AI assistants in 2025, 87% of it ChatGPT. Conductor, 1,215 domains, 3.3bn sessions. [C]

Seer Interactive's 2026 update (53 brands, 2.43 billion organic impressions) gives the most useful shape: organic CTR of 3.35 percent where no AI Overview appears, 2.07 percent where one appears and the brand is cited, 0.94 percent where one appears and the brand is not cited. [c] Amsive's 700,000-keyword study found non-branded CTR down 20 percent and branded CTR up 19 percent where AI Overviews appear, because they rarely appear on branded queries. [c] Seer also measured AI Overviews on 36 percent of informational queries, 8 percent of commercial and 5 percent of transactional. The commercial-intent pages a consultancy lives on are the least exposed part of the results page.

AI referral traffic is small and growing fast. The headline conversion multiples ("23x", "4.4x", "converts 42 percent better") rest on a single site, an undisclosed method, or a baseline that flipped sign inside twelve months (Adobe's US retail data went from 38 percent worse to 42 percent better between March 2025 and March 2026). [d] Treat AI traffic as directionally high-intent and unquantifiable.

What the platforms now say and provide

- **Google, 15 May 2026.** A new guide, "Optimizing your website for generative AI features on Google Search", states that generative features "are rooted in our core Search ranking and quality systems", that lms.txt and similar files are unnecessary and "Google Search ignores them", that chunking content and keyword variants are unnecessary, and that structured data is "not required for generative AI search". It asks for "unique expert or experienced takes that go beyond common knowledge" and names generic "7 tips" explainers as the content to avoid. [A]

"UNIQUE EXPERT OR EXPERIENCED TAKES THAT GO BEYOND COMMON KNOWLEDGE."

GOOGLE, OPTIMIZING YOUR WEBSITE FOR GENERATIVE AI FEATURES, 15 MAY 2026

- **Google Search Console, 3 June 2026, global by 31 August.** A Search Generative AI performance report showing impressions in AI Overviews and AI Mode by page, country, device and date. No clicks, no CTR, no queries. Alongside it, a property-level control to exclude a site from generative features entirely, with no effect on other rankings. [A]

- **Google Analytics 4, 13 May 2026.** A native "AI Assistant" default channel keyed on a referrer list Google does not publish. Copilot sends no referrer and mobile apps strip it, so a material share of AI traffic still lands in Direct. [\[A\]](#)
- **Microsoft Bing, 26 to 27 February 2026.** Rewrote its Webmaster Guidelines to cover Copilot and grounding, defined GEO as "content eligibility for grounding and reference in AI responses" with no guarantee of citation, made NOARCHIVE and NOCACHE govern Copilot use, softened its AI-content stance to target "large-scale content generated without oversight", and added abuse categories for language engineered to trigger citations and for prompt injection. [\[A\]](#) (page is JavaScript-only; wording via two trade reports)
- **Bing Webmaster Tools, 10 February 2026.** An AI Performance report with total citations, cited pages and "grounding queries" (the reformulated queries Copilot used when retrieving). It is the only first-party citation-level data any engine provides. [\[A\]](#)
- **Structured data.** HowTo rich results ended in 2023; seven more types were dropped in June 2025; FAQ rich results ended for all sites in May 2026 with API support removed in August 2026. Google has never stated that schema helps AI citation. Microsoft's Fabrice Canel has said schema helps its language models understand content. [\[A\]](#)
- **The rater guidelines are dated 11 September 2025.** A "June 2026 update" adding "Synthetic Authority" and "Verifiable Real World Expertise" sections circulates on SEO blogs. The section headings quoted do not match the live PDF. Treat as fabricated. [\[U\]](#)

PART I • SECTION 3

WHAT MOVES ORGANIC RANKINGS

Two primary documents changed what can be said with confidence: the May 2024 leak of Google's Content Warehouse API reference (2,596 modules, 14,014 attributes) and Judge Mehta's August 2024 liability opinion in *US v. Google*, with its sworn testimony from Pandu Nayak and Eric Lehman. Neither reveals weights. Both reveal what Google stores and, in the court's case, what it relies on.

ESTABLISHED

User interaction. Finding of fact 96: Navboost "pairs queries and documents through memorizing user click data" over 13 months. Finding 102: "the more recent LLM signals did not replace Navboost". Finding 103: Navboost "can also beat out LLMs" on freshness. The leak names the attributes (goodClicks, lastLongestClicks, chromeInTotal). The practical target is the last long click: the visit after which the searcher does not return to the results page. [A]

"THE MORE RECENT LLM SIGNALS DID NOT REPLACE NAVBOOST."

US V. GOOGLE, FINDING OF FACT 102, AUGUST 2024

ESTABLISHED

A site-level quality score exists. The trial record describes Q* as a largely static, query-independent site quality score; the leak has siteAuthority, siteFocusScore and siteRadius (how far a page's topic sits from the site's centre) and hostAge used "to sandbox fresh spam". Thin or off-topic pages are a site-level liability, not a neutral addition. [A]

ESTABLISHED

Google removes thin pages regardless of site size. Its crawl-budget guidance applies to sites over roughly 10,000 pages; below that, "Crawled, currently not indexed" is a quality verdict, not a budget problem. A May 2025 index purge reportedly removed basic FAQ and definition pages and pages "that previously relied on internal linking for visibility" first. [A] [D]

ESTABLISHED

Links count, but less than they did. Gary Illyes, 2023: not in the top three. Google's spam policy lost the word "important" in March 2024. Link value is tiered by the quality of the linking page. Population-scale correlations between referring domains and traffic are strong but contaminated by reverse causation. [A] [C]

ESTABLISHED

For sector-qualified B2B terms, links are not the binding constraint. UK results pulled today via the Ahrefs API: for "b2b pr agency", pages with zero referring domains sit at positions 4, 7, 8 and 10 (the worked-example site in section 9 is the one at 4). For "healthcare pr agency", positions 4 to 6 are domains rated DR 9 to 14. For generic head terms ("seo agency") the top ten carries hundreds to thousands of referring domains and directories take two of eight organic slots. [D] (four SERPs, indicative)

ESTABLISHED

Titles are rewritten most of the time. 61.6 percent across 80,959 titles (Zyppy), 76 percent across 30,000 keywords (McAlpin, 2025). Titles of 51 to 60 characters have the lowest rewrite rate; over 70 characters, 99.9 percent. Brand names were removed in 63 percent of rewrites; pipes are replaced twice as often as dashes. [C]

ESTABLISHED

Internal links and anchor variety correlate with clicks. Zyppy, 23 million internal links: pages with 40 to 44 inbound internal links earned about four times the Search Console clicks of pages with 0 to 4, with decline past roughly 50; anchor-text variety was the strongest relationship in the data. [C]

ESTABLISHED

Core Web Vitals are a confirmed, minor signal. Google: "not giant factors in ranking". Ahrefs, 42 million pages: not significant. Pass the thresholds and stop. [A] [C]

CORRELATIONAL

Brand demand predicts visibility. Across 75,000 brands, Spearman correlations with AI Overview presence were 0.664 for web mentions, 0.392 for branded search volume, 0.326 for Domain Rating and 0.218 for backlinks. No study shows that raising brand demand causes non-brand rank gains; both may flow from being a known brand. Fishkin's reading of the leak: for most small businesses "SEO is likely to show poor returns until credibility and navigational demand establish". [\[C\]](#) [\[D\]](#)

FOLKLORE

Content length as a factor (Google: no preferred word count; Backlinko, 11.8 million results: no direct correlation). Publishing frequency (no primary support). Link velocity (the only documented velocity attribute is a spam detector). Domain Authority as Google's own measure (Google: third-party scores "don't correspond to any of Google's own signals"). Meta keywords, readability scores, exact word counts. [\[A\]](#)

PART I · SECTION 4

WHAT MOVES AI CITATION

The pipeline, and where each study sits

The single most useful framing comes from Olivier Martinez's July 2026 survey of 45 GEO studies (Sciences Po, no commercial interest; the widely shared Limy blog post is a popularisation of it). [A] Citation is the end of a pipeline: does the engine search at all, is the page crawled and indexed, is it retrieved, does it survive reranking into the context window, is it cited, is its content actually absorbed into the answer, does anyone click. Almost every GEO claim measures stage five with stages one to four held fixed. Google describes the front of that pipeline as "query fan-out": the engine issues many related sub-queries and retrieves passages for each. [A]

What holds up

- **Relevance to the literal question.** The dominant determinant of whether a retrieved source is used, in the academic record (Wan et al., ACL 2024; Zhang et al. 2026, $r = 0.43$ between judged relevance and influence on the answer). In ChatGPT data, heading-to-query similarity above 0.90 lifted citation rate to about 41 percent from about 30 percent, the strongest page-level signal AirOps found across 353,799 pages. [A] [C]
- **Extractable evidence.** The 2024 GEO paper's best-performing edits were adding quotations (+41 percent position-weighted share), statistics (+30 to 33) and cited sources (+28); keyword stuffing performed below baseline. Zhang et al. found uplift for numerical statistics (+62), definitions (+57) and comparisons (+55), and a slight penalty for Q&A format (-5.7). Cited passages carry roughly three times the proper-noun density of typical text. [A] [C]
- **Position within the page.** 44.2 percent of ChatGPT citations came from the first 30 percent of a page, and 78.4 percent of heading-linked citations sat under H2s (Indig, 18,012 verified citations). This is a 1.5x over-representation, not a rule that the rest is ignored. [C]
- **Focus over breadth.** Pages covering a quarter to a half of ChatGPT's fan-out sub-queries were cited more than pages covering all of them, after controlling for relevance. Long pages are fine when every section answers a distinct sub-question; padding around one answer is not. [C]
- **Freshness, for time-sensitive topics.** 65 percent of ChatGPT crawler hits land on content under a year old (Seer, log data); AI-cited pages average 25.7 percent younger than organic results (Ahrefs, 17 million citations). But among pages ChatGPT retrieved, the ones it cited were older (median about 500 days). The likely reconciliation is that recency helps retrieval and ordering while age and corroboration help selection. Unresolved. [C]
- **Presence in the sources engines already cite.** Earned media took 84 percent of ChatGPT, Claude and Gemini citations across 25 million links (Muck Rack). Owned domains took 47.5 percent overall but 59.8 percent inside Google AI Overviews and 28.9 percent on Perplexity (Otterly). YouTube mentions are the strongest single correlate of AI visibility on all three engines Ahrefs tested (0.71 to 0.74). [C]
- **Server-rendered HTML.** None of the OpenAI, Anthropic, Perplexity or Meta crawlers executes JavaScript; only Googlebot and Applebot render (Vercel and MERJ, 1.3 billion fetches). Client-rendered content is invisible to ChatGPT, Claude and Perplexity. [B]

What is contradicted or unsupported

CLAIM	WHAT THE EVIDENCE SAYS	GRADE
GEO rewrites deliver 30 to 40 percent gains	Only inside a fixed five-document lab context (GEO paper). C-SEO Bench: 3 of 54 method and domain combinations positive. SAGEO Arena: body rewrites cut retrieval presence 9 percent, post-rerank top ten 16 percent, citations 6 percent.	[A]
Schema markup improves AI citation	Ahrefs' matched difference-in-differences (1,885 pages that added JSON-LD vs 4,000 controls): AI Overviews -4.6 percent (p about 0.0004), AI Mode and ChatGPT not significant. Google: no special schema needed.	[B]
llms.txt helps	Across 137,210 domains, 97 percent of llms.txt files received zero requests in May 2026; retrieval bots were 1.1 percent of the fetches that did occur. A 900-domain, 191-day log study saw no frontier AI bot request it once. Google: ignored. Mueller: "none of the AI systems use it".	[B] [A]
Markdown mirrors or "clean HTML for AI"	No engine documents consuming them; no outcome evidence exists.	[U]
FAQ blocks as a universal tactic	Slightly negative in the one large feature study; FAQ rich results are gone. Question-form H2 with a direct answer beneath is the format that works; a bolted-on FAQ that repeats the body competes with it.	[C]
Domain authority drives citation	Surfer, 5 million citation URLs: PageRank, harmonic centrality and Domain Score correlate near zero or slightly negatively. Ahrefs: DR 0.27 to 0.33.	[C]
"Authoritative tone" rewrites	+12 percent in the GEO paper; "weak and unstable" in the survey; possibly counterproductive.	[A]
Gartner: search volume down 25 percent by 2026	Google reports query growth and AI Mode at a billion monthly users. The decline is in clicks, not queries.	[A]
"Wikipedia is 47.9 percent of ChatGPT citations"	It is 47.9 percent of the top-ten source share in one dataset; 7.8 percent of all ChatGPT citations in the same dataset.	[C]
Digital PR causes AI citations	No controlled before-and-after study exists. Every brand-mention result is cross-sectional correlation among brands that are already large.	[U]

Fragmentation and volatility

Any two engines share about 17 percent of cited domains; only 3.8 percent of domains were cited by all four of ChatGPT, Perplexity, Gemini and AI Overviews for the same prompts (Writesonic, 161,286 prompts). ChatGPT is the structural outlier, leaning on Wikipedia and editorial outlets; Google's surfaces barely cite Wikipedia and lean on YouTube, Reddit and LinkedIn; Perplexity behaves most like a conventional search engine (91 percent domain overlap with Google's top ten). [C] Day to day, 69 percent of an answer's sources change (Gemini 88 percent, ChatGPT 79, AI Mode 76, Perplexity 44), and ChatGPT skipped web search

entirely on 57.8 percent of runs in one academic sample.

[C] [A] The overlap between AI Overview citations and Google's own top ten fell from 76 percent (2025, partly a parsing artefact) to 17 to 38 percent across four 2026 datasets, which Google attributes to fan-out. Ranking somewhere on the topic from a domain Google trusts matters more than ranking first for the exact phrase. [C]

IMPLICATION

A single "GEO strategy" is not a coherent object. What generalises is: be retrievable (indexed, server-rendered, on-topic), be quotable (a self-contained answer with a sourced number under a matching heading), and be present in the third-party sources each engine already trusts. Measure per engine, over repeated runs, and report ranges.

02

PART TWO • SECTIONS 5 AND 6

THE AUDIENCE AND THE PAGE

What buyers and journalists say they respond to, triangulated across publishers with different interests, and the article template that follows from it.

COVERS

Buyer and journalist surveys, measurement reality, the 14-element Insight template, writing rules

KEY THEME

Original research, named experts and a position the sector will argue with

PART II • SECTION 5

WHAT BUYERS AND JOURNALISTS RESPOND TO

Every source in this field sells something adjacent to what it recommends, so the numbers below are triangulated across publishers with different interests and graded accordingly.

Buyers

- **Thought leadership reaches out-of-market buyers and they act on it.** Edelman and LinkedIn 2024 (about 3,500 decision-makers, seven countries): 75 percent researched a product or service they were not considering because of a piece of thought leadership; 86 percent would be more likely to invite a consistent publisher into an RFP; 60 percent would pay a premium; 70 percent of the C-suite have questioned an incumbent supplier because of it. Fewer than half rate what they read as good; 15 percent as very good. [c]
- **Hidden buyers want to be challenged.** Edelman and LinkedIn 2025 (about 2,000): 71 percent of the people who influence a decision have little or no contact with sales; 86 percent do not want thought leadership to validate their thinking; 91 percent value insight that surfaces risks they had not recognised; 57 percent prefer quick takeaways to academic depth; 65 percent prefer a human, less formal tone; 53 percent say strong thought leadership makes brand recognition matter less. [c] No 2026 edition exists, and the widely quoted 2025 finding that buyers penalise "AI-sounding" content could not be found in the report. [u]
- **The preferred vendor is usually chosen before first contact.** 6sense 2025 (4,000+ buyers): 94 percent of buying groups ranked a preferred vendor before contacting anyone and bought it 77 percent of the time; first contact now comes 61 percent of the way through the journey. [c]

- **Named technical experts beat executives on trust.** Edelman Trust Barometer 2026 (33,938 respondents): a company's in-house technical experts are on average 22 points more credible than its CEO. Scientists and "someone like me" tied as the most trusted sources at 74 percent in 2024. Political figures sit near the bottom for the general public. [c]
- **High-growth professional services firms behave differently.** Hinge 2025 (770 firms) and 2026 (495): high-growth firms grow about four times as fast, spend 12 percent of revenue on marketing against 5 percent for no-growth firms, are 2.5 times more likely to put subject-matter experts forward, track SEO and GEO at 50 percent against 36 percent, and in 2026 shifted their top priority from creating content to promoting it and building individual thought leaders. [c] Hinge's older referral research: 81.5 percent of firms have received a referral from a non-client, seeded by speaking (30 percent), articles (20) and social (17); 52 percent of referred firms are ruled out on a weak website before a call. [c]

"52 PERCENT OF REFERRED FIRMS ARE RULED OUT ON A WEAK WEBSITE BEFORE A CALL."

HINGE RESEARCH INSTITUTE, REFERRAL RESEARCH

- **Research reports are the under-supplied format.** CMI 2025 (980 B2B marketers): research reports rank joint third for results (45 percent) with far lower adoption than articles or video. CMI 2026: LinkedIn (76 percent), email newsletters (54) and speaking (52) are the thought-leadership channels rated effective; 56 percent cannot attribute ROI. [c]

Journalists

- Cision 2025 (3,000+ journalists): 55 percent used original research or data from communications partners in their work; 86 percent immediately reject an off-beat pitch. Cision 2026 (1,899): 66 percent rely on PR-supplied material for story ideas, the leading source; journalists want data, embargoes and access to experts; 53 percent object to AI-written pitches. [\[c\]](#)
- Muck Rack 2026 (897 journalists): 86 percent say at least some stories start with a pitch; preferred pitch is a one-to-one email under 200 words, before noon, with one follow-up; LinkedIn is the platform they trust most (58 percent). [\[c\]](#)
- Links follow original research. 94 percent of blog posts earn zero external links (Backlinko, 912 million posts); annual survey reports earned roughly ten times the linking domains of comparable content (BuzzSumo and Majestic); Orbit Media's annual survey earned 430 links in six months; Ahrefs earned 36 editorial links from 515 pitches for a statistics page. [\[c\]](#) [\[D\]](#)

Measurement reality

Refine Labs found a 90 percent gap between software attribution and self-reported attribution; its podcast drove 53 percent of self-reported revenue and 0 percent of software-attributed revenue. Dreamdata's B2B journeys average 211 days and 76 touches. The realistic horizon for a new content programme is six to twelve months to first attributable enquiries, and two to three annual cycles before a recurring index becomes a citation other people reach for. The documented failure mode is stopping at month four because two-week metrics look flat. [\[D\]](#) [\[c\]](#)

PART II • SECTION 6

HOW TO STRUCTURE AND WRITE AN INSIGHT

Everything in this section is tied to evidence in sections 3 to 5. Elements marked (inference) are reasonable conclusions rather than directly tested.

The page, in order

- 1. Title, 50 to 60 characters.** A natural-language statement of the specific question or claim, mirroring the H1, dash rather than pipe, brand last or omitted on deep pages. Google keeps numbers 97 percent of the time when title and H1 agree. Title-to-prompt similarity predicts ChatGPT citation (0.60 vs 0.48). [\[C\]](#)
- 2. H1 as the question or claim in plain words.** Google builds title links from headings when it rewrites; heading-query match is the strongest page-level ChatGPT signal. [\[A\]](#) [\[C\]](#)
- 3. Byline with credentials, published date, updated date only when substance changed, and for health or financial (YMYL) content a named clinical or professional reviewer.** Google's self-assessment asks whether it is "self-evident who authored your content"; the rater guidelines make "who is responsible" central and rate fake or inflated author profiles Lowest. No causal ranking test of bylines exists; the downside of omitting them on YMYL content is documented. [\[A\]](#)
- 4. Opening block, first 100 to 150 words.** The direct answer in two or three sentences using definitional phrasing ("X is", "X refers to"), one concrete number with its source, one sentence on why it matters to this reader. Cited passages use definitional phrasing twice as often as uncited ones; 44 percent of citations come from the first 30 percent of a page; Nielsen Norman's eyetracking puts the most important points in the first two paragraphs; featured-snippet paragraphs run 40 to 55 words. [\[C\]](#) [\[B\]](#)
- 5. Optional key-findings box, three to five bullets.** Each a complete, specific, sourced claim, not a teaser. (inference from 4)
- 6. Body sections.** Each H2 phrased as the sub-question a fan-out query would use; first sentence answers it; two or three short paragraphs; a table only where the content is a comparison, a numbered list only where it is a procedure; named sources inline. Google scores passages and has said heading order and count do not matter to ranking, so use order for readers and wording for retrieval. 78 percent of heading-linked citations sit under H2s. [\[A\]](#) [\[C\]](#)
- 7. Scope discipline.** Cover the sub-questions that belong to this URL's question, not every adjacent one. Pages covering 26 to 50 percent of fan-out sub-queries out-cite exhaustive pages; Google has no preferred word count. [\[C\]](#) [\[A\]](#)
- 8. Original material.** Your data with a method line, your named quotes, your position. This is what Google's self-assessment calls "original information, reporting, research, or analysis" and "beyond the obvious"; it is what the GEO paper's three best edits add; it is what 86 to 91 percent of hidden buyers say they want. [\[A\]](#) [\[C\]](#)
- 9. Sources and method block.** Every statistic with source, date and link; what was measured and how. Eleven percent of AI Overview claims are unsupported by their cited pages; about 16 percent of engine citations are themselves AI-generated pages; fabricated claims dressed as corroborated were absorbed by some models 78 to 92 percent of the time. An unsourced "studies show" is now a reputational liability, because an engine will repeat it with your name attached. [\[A\]](#) [\[C\]](#)
- 10. Author page linked from the byline,** with bio, credentials, other work and profile links; Article and Person or ProfilePage JSON-LD whose names and dates match the visible page. Not a ranking lever; a trust and disambiguation asset. [\[A\]](#)
- 11. Related-questions links** to sibling Insights rather than an FAQ block that repeats the body. (inference; FAQ rich results are gone and duplicate passages compete for the same sub-query)

12. **A change-log line on refresh** ("Updated 2 September 2026: replaced 2024 ONS figures with the 2025 release"). Google lists date-bumping without change as a bad practice; Mueller calls it "just noise". 76 percent of ChatGPT's most-cited pages had been updated within 30 days. [\[A\]](#) [\[C\]](#)
13. **Charts and images with descriptive alt text, short video where useful.** Google's guide explicitly recommends images and video; evidence of citation lift is vendor-only. [\[A\]](#) [\[D\]](#)
14. **UK signals.** UK institutions and datasets (NHS, CQC, ONS, DHSC), sterling, UK dates, UK spelling. Spelling is not an evidenced ranking factor; local entities and sources are relevance signals for UK-intent sub-queries and trust signals for UK readers. [\[A\]](#)

Writing rules that follow from the evidence

- **Lead with the claim, then the evidence, then the consequence.** Buyers want quick takeaways (57 percent) and the engines lift the first third of the page.
- **Take a position the sector will argue with, and back it with a number you own.** Stated preference (86 percent) and link behaviour both favour opinionated pieces; no controlled study shows contrarian content converts better, so pair the position with data rather than replacing it.
- **Write at the level of an informed peer.** Cited ChatGPT passages read at roughly grade 16, simpler than the grade 19 material around them but nowhere near grade 8. Readability scores do not correlate with rank (Portent, 5.8 million pages).
- **Use proper nouns.** Name the regulator, the report, the person, the date. Entity density is what distinguishes cited passages.
- **Cut anything the model already knows.** Generic explainers are the content Google's guide names as low value. If a paragraph could appear on a competitor's site unchanged, it is not earning its place.

- **AI assistance is permitted; unedited AI output at scale is not.** The rater guidelines judge effort, originality and added value "no matter how" content is created, and rate fake author profiles Lowest. Orbit Media's survey: the 10 percent who let AI write whole articles report the weakest results. Disclose material AI assistance; disclosed-then-discovered is the worst case in the NIM trust experiments. [\[A\]](#) [\[C\]](#)

Practices to drop

- Writing to a word count or padding to match the top result.
- Keyword-form rewriting of headings and sentences. The only GEO method that performed below baseline; Google says its systems handle synonyms.
- Bolting a generic FAQ block onto a page that already answers the question.
- Bumping the visible date or dateModified without substantive change.
- Unsourced statistics and round numbers from memory.
- Bracketed title suffixes and pipes (rewritten 78 and 41 percent of the time respectively).
- "Ultimate guide" pages that cover every adjacent question on one URL.
- Optimising for one engine's citation style. Only 2.4 percent of URLs are cited by ChatGPT, Perplexity and AI Overviews for the same prompt.
- Treating "information gain score" as a live ranking factor. The patent (US11354342B2) is scoped to assistant re-ranking after documents have been seen.

03

PART THREE · SECTIONS 7 AND 8

PLUMBING AND MEASUREMENT

The eligibility gate, ordered by evidence strength and effort, and a measurement stack that separates census sources from sampled ones.

COVERS

Rendering, crawler access, firewall and robots.txt, sitemaps, Bing, structured data, canonicals, reporting rules

KEY THEME

Technical work is where visibility is lost, not where it is won

PART III • SECTION 7

TECHNICAL PLUMBING

This layer is the eligibility gate, not where visibility is won. Ordered by evidence strength and effort.

1. **Rendering.** Every Insight, author page and service page must be complete in the initial HTML. Non-Google AI crawlers never execute JavaScript. Test with `curl -A "GPTBot"` and confirm body text, headings, links and JSON-LD are present; no content behind client-only tabs or expanders. [\[B\]](#)
2. **Crawler access.** Three classes exist and the consequence of blocking each differs. Training crawlers (GPTBot, ClaudeBot, the Google-Extended token, CCBot): blocking excludes content from future training, not from answers. Retrieval indexers (OAI-SearchBot, Claude-SearchBot, PerplexityBot, Googlebot, Bingbot): blocking is self-exclusion from citations; OpenAI's documentation says opted-out sites "will not be shown in ChatGPT search answers". User-triggered fetchers (ChatGPT-User, Claude-User, Perplexity-User): mostly ignore robots.txt; block at WAF level only if you mean it. Google-Extended does not affect Search, AI Overviews or AI Mode. [\[A\]](#)
3. **Vercel firewall.** The AI Bots managed ruleset is a single Allow, Log or Deny switch with no training-versus-retrieval split; Deny removes OAI-SearchBot and PerplexityBot along with GPTBot. Leave it on Log for free crawl telemetry. Do not run Attack Mode as a standing setting. Bot Protection needs a Bypass rule for the retrieval and user fetchers (by user agent, and by published IP range where possible); Vercel supports Web Bot Auth signatures as a match condition. Vercel states Bot Protection does not work behind a Cloudflare proxy. [\[A\]](#)
4. **Cloudflare, if it proxies the domain.** New zones block AI crawlers by default since July 2025, and from 15 September 2026 Cloudflare will block "mixed-use" crawlers on ad-bearing pages by default for new and free-tier zones. If your domain is proxied (orange cloud) rather than DNS-only, Cloudflare's AI Crawl Control decides which bots reach the origin regardless of Vercel's rules. Check. [\[A\]](#)
5. **robots.txt.** Explicitly allow the retrieval indexers; make the training-crawler decision consciously (default allow for a firm that wants models to know who it is; Anthropic frames blocking ClaudeBot as exclusion from training only); never disallow `/_next/static/`; include the sitemap URL. Add Content-Signal: `search=yes, ai-input=yes` (Cloudflare's CC0 vocabulary, interim ahead of the IETF AIPREF drafts). [\[A\]](#)
6. **Sitemap with honest lastmod** sourced from the CMS revision date, never `new Date()` at build time, which discredits every entry. Google ignores priority and changefreq. Submit an Atom feed for Insights as a supplementary sitemap. [\[A\]](#)
7. **Bing Webmaster Tools and IndexNow.** Verify the site, enable the AI Performance report, and POST changed URLs to `api.indexnow.org` on publish and update from Payload. Microsoft says 18 to 22 percent of newly clicked URLs arrive via IndexNow; there is no independent test, but for a site Bing crawls infrequently it is the cheapest route into the index Copilot grounds on, and ChatGPT still blends Bing's index. Google does not participate. [\[A\]](#)
8. **Structured data, server-rendered.** Organization on every page with `sameAs` to LinkedIn, Companies House and social profiles; Article with a Person author whose `url` points to an internal ProfilePage; ProfilePage on author pages; BreadcrumbList. Keep it consistent with visible content. It is hygiene for rich results and entity disambiguation, not a citation lever. A Wikidata item for a non-notable firm fails the notability test and has no measured effect. [\[A\]](#) [\[B\]](#)
9. **Canonical hygiene.** One host with an edge 301, absolute self-canonical on every route, canonical strips unknown parameters, parameter spam returns 200-with-canonical or 301, never 429 or 5xx. Google treats every 4xx except 429 identically; 410 has no documented speed advantage over 404. Noindex only on crawlable paths. [\[A\]](#)

10. Core Web Vitals. Pass LCP, INP and CLS at the 75th percentile in CrUX, then stop. AI crawlers do not execute JavaScript and never experience INP. [\[A\]](#)

NOT WORTH DOING, WITH EVIDENCE

llms.txt beyond the file that exists (zero frontier-bot requests in two independent log studies; leave it, do not extend it). Markdown alternates. FAQ or HowTo schema. Special "AI schema". Sitelinks search-box markup (feature removed October 2024). 410 instead of 404 for spam URLs. Blocking training crawlers by reflex. Vercel AI Bots on Deny. Buying links, or paying for directory placements that are anything other than a real listing with real reviews. Chasing CWV beyond the thresholds.

PART III • SECTION 8

MEASUREMENT

Separate census sources (complete but narrow) from sampled sources (broad but noisy), report ranges not points, and always show the chain from crawl to citation

to referral so that "AI visibility" is anchored in something observable.

LAYER	SOURCE	CADENCE	WHAT IT CAN AND CANNOT TELL YOU
Access	Vercel Firewall traffic (AI Bots ruleset on Log)	Weekly	Requests per bot, allowed vs challenged, time from publish to first fetch by OAI-SearchBot, PerplexityBot, Bingbot, Googlebot, 404 share per bot. Census. Cannot tell you whether a fetch became a citation.
Inclusion	Search Console Generative AI report; Bing Webmaster Tools AI Performance; URL Inspection	Monthly	AI Overview and AI Mode impressions by page, country and device (no clicks, no queries). Copilot citations, cited pages and grounding queries (no clicks, no API). Index coverage. Census for two engines only.
Citation	Fixed prompt set, 25 to 40 prompts, frozen per quarter; run on ChatGPT (search on), Perplexity, AI Mode, Gemini, Copilot, Claude, logged out, UK location	Weekly or fortnightly	Appearance rate per engine with a binomial interval. Minimum four runs per prompt per engine before comparing periods; target 100+ prompt-runs per engine per period. Ahrefs Brand Radar or similar as a second opinion for breadth, not the KPI; one independent test found it undercounting 40-fold on niche prompts. Sample.
Traffic and outcome	GA4 AI Assistant channel plus a custom "AI" channel group keyed on referrer; PostHog referring domain and utm_source=chatgpt.com; "How did you hear about us" free text on every enquiry form	Monthly	Sessions, engaged sessions, enquiries. Copilot and Claude send no referrer and mobile apps strip it, so watch Direct for anomalies. Self-reported attribution is the only thing that catches dark-social and podcast influence (Refine Labs' 90 percent gap).

Reporting rules. Every AI metric carries its source class and its interval. No month-on-month claim on the citation layer without 100 prompt-runs per engine in each month. Never present tool "mentions" as traffic or Search Console AI impressions as clicks. Annotate engine changes (model releases, the September 2025 Reddit-share collapse in ChatGPT, the January 2026 Gemini 3 upgrade that replaced 42 percent of AI Overview domains) so volatility is not misread as performance. Segment Search Console CTR by branded and non-branded, page type, country and device before characterising it at all.

04

PART FOUR • SECTIONS 9 AND 10

ONE SITE, ONE PROGRAMME

The evidence applied to a real UK B2B agency website, pulled live on 2 September 2026 and anonymised, and the three-tier programme that follows from it.

COVERS

Search Console clicks, snippet substitution, source-level truncation, three competitors, Tiers 1 to 3

KEY THEME

Snippet integrity and content were the binding constraints, not authority

PART IV • SECTION 9

A WORKED EXAMPLE: ONE SITE TODAY

A real UK B2B PR and public affairs agency site, pulled live on 2 September 2026 and anonymised here. Figures are stated; causes are hypotheses unless marked otherwise.

Where the clicks come from

Search Console, 1 June to 31 August 2026. Almost every click is branded or navigational: the agency's own name accounted for 215 clicks, a flagship named tool on the site a further 70, team members' names a further 30 or so. Commercial service pages collect tens of thousands of impressions and almost no clicks.

PAGE	CLICKS	IMPRESSIONS	CTR %	AVG POSITION
Homepage	143	39,603	0.36	52.7
Sector page A	1	32,363	0.003	53.2
Sector page B	6	21,474	0.03	49.6
Service page A	4	18,422	0.02	10.2
Service page B	6	12,132	0.05	20.6
Sector page C	2	10,883	0.02	40.8
Sector page D	3	10,075	0.03	31.6
Service page C	7	7,538	0.09	12.8
Vendor-directory page	2	7,070	0.03	6.4

Several commercial queries sit inside the top ten with click rates far below any published curve: "b2b pr agency london" at average position 3.9 with 1,463 impressions and 1 click; "b2b pr agency" at 7.0 with 4,119 impressions and 1 click; "b2b pr agencies" at 4.2 with 915 impressions and 1 click. Those average positions blend countries and devices, and ads, a local pack or an AI Overview may sit above the first organic result, so the numbers are not evidence of a broken page on their own. Two things checked today are evidence.

Finding one: Google is not using the site's snippets

In a non-personalised UK search for its primary commercial term today, the site's page is the first organic result. The description Google shows is a sector-flavoured sentence plus a full postal address, lifted verbatim from the sitewide Organization JSON-LD description block. It is site-wide boilerplate, and for a B2B buyer the sector-specific framing answers a question they did not ask. The page's own meta description, which does address B2B buyers, is being ignored. [A] (observed)

Finding two: titles and descriptions are truncated in source

Fifteen key pages were fetched as Googlebot. Eleven have a literal ellipsis character cutting the title or meta description mid-claim, in three cases mid-brand-name. The cause is in the site's metadata-generation helper: a function that cuts at a word boundary and appends an ellipsis, applied at 60 characters to titles and 155 to descriptions. Its comment says the budgets exist because "Ahrefs flags beyond these". It was introduced to clear an Ahrefs site-audit warning and extended in two later changes. [A] (verified in source)

“THE SITE IS HARD-TRUNCATING AUTHORED COPY TO SATISFY A THIRD-PARTY AUDIT TOOL AND HANDING GOOGLE A VISIBLY UNFINISHED SENTENCE.”

SECTION 9, FINDING TWO

Google's documentation says there is no length limit on titles or descriptions, that display truncation is Google's job, and that incomplete or boilerplate descriptions are a reason it substitutes its own snippet. The site is hard-truncating authored copy to satisfy a third-party audit tool and handing Google a visibly unfinished sentence; on the highest-value page Google has responded by falling back to schema text. Whether this explains the whole click-rate gap is unproven. It is the cheapest testable intervention on the site: change the policy so the helper never emits an ellipsis into a title or description (keep the budgets as an editorial guideline), rewrite the eleven affected strings to be complete and intent-matched within budget, give the Organization description a neutral sentence that reads acceptably as a fallback snippet, and read Search Console CTR on the affected pages at positions 1 to 10 over four weeks against the June to August baseline.

Authority, content and citations against three tracked competitors

SITE	DR	REF. DOMAINS (FOLLOWED)	GB ORGANIC VISITS/MO (EST.)	GB KEYWORDS	AI CITATIONS (PAGES)	AI MODE + AIO	COPILLOT
This site	15	663 (190)	527	56	51 (30)	6	0
Competitor A	57	1,267 (540)	2,522	292	40 (22)	6	1
Competitor B	46	1,196 (506)	4,221	314	259 (89)	60	4
Competitor C	56	1,582 (961)	711	45	3 (3)	0	0

Ahrefs, 2 September 2026, country GB, subdomains mode. Competitors anonymised. AI citation counts are from Ahrefs' sampled prompt corpus and measure presence in that sample, not true share.

Competitor C has nearly four times this site's Domain Rating and five times its followed referring domains, yet only about a third more UK organic traffic and fewer AI citations, because it publishes little indexable content. Competitor B has the most content and the most AI citations, concentrated in Google AI Mode and AI Overviews, exactly the surfaces Otterly finds most receptive to owned content. This site's 51 citations exceed Competitor A's despite a Domain Rating of 15

against 57. This is the section 3 evidence in miniature: authority without content produces neither rankings nor citations; content is the binding constraint, then earned mentions. This site's zero on Copilot is consistent with Bing seeing it rarely, which the Bing Webmaster Tools and IndexNow work in section 10 addresses directly.

Other standing findings, re-read against the evidence

- **One service line** remains the largest unserved commercial cluster (about 30,000 impressions over six months at positions 24 to 37, keyword difficulty 1 to 14). Its hub page shipped recently; the evidence says give it 20 to 45 contextual internal links with varied anchors, a sourced number in its first paragraph, and three to six months. [c]
- **Keyword cannibalisation.** One near-zero-difficulty commercial term is split across three overlapping pages on the same topic. The Ahrefs evidence on cannibalisation is that consolidation only helps when intent is identical; here it is, so the retitle-as-hub decision should be re-read in the next Search Console cycle and escalated to a redirect if the split persists.
- **A family of programmatic vendor-directory pages** collects large impression volumes on vendor brand names (one single vendor query drew 5,137 impressions, 1 click) where the searcher wants the vendor's own site. These are not harmful in themselves, but the trial record's site-level quality score and the leak's siteRadius make the standard clear: each page must carry information that is not available elsewhere, or it is a dilution risk. Batch cautiously and watch "Crawled, currently not indexed" as a quality verdict.
- **A prior position degradation** (average position roughly 20 to 38 across the GB market, unrecovered at last check) should be read as its own line with its own evidence, separate from any concurrent spam campaign or site migration.

PART IV • SECTION 10

THE PROGRAMME

Priority is set by evidence strength times expected effect divided by effort. Each item names how it will be measured.

Tier 1: this month, one developer session each

1. **Snippet integrity.** Remove the ellipsis truncation from titles and descriptions; rewrite the eleven affected strings complete and intent-matched; neutralise the Organization description; 50 to 60 character titles, dash separator, brand last. Measure: CTR on pages at positions 1 to 10, four weeks vs the June to August baseline, segmented by country and device. [\[A\]](#)
2. **Crawler access audit.** Confirm the Vercel AI Bots ruleset is on Log, Attack Mode is off, the existing bypass rule covers OAI-SearchBot, ChatGPT-User, Claude-SearchBot, Claude-User, PerplexityBot, Perplexity-User, Bingbot and Applebot, and that Firewall traffic shows them allowed rather than challenged. Confirm whether Cloudflare proxies the apex. Add the Content-Signal line to robots.txt. Measure: bot request counts per week, time from publish to first OAI-SearchBot and Bingbot fetch. [\[A\]](#)
3. **Bing.** Verify in Bing Webmaster Tools, turn on AI Performance, wire IndexNow to the CMS's publish and update hooks. Measure: Copilot citations from a baseline of zero. [\[A\]](#)
4. **Measurement plumbing.** GA4 custom AI channel group above Referral; an analytics referrer filter; "How did you hear about us" free text on every enquiry route; Search Console Generative AI report read monthly; a 30-prompt UK set (brand, category, "who to hire" and the site's core service lines) frozen per quarter and run four times per engine per month. [\[A\]](#)
[\[C\]](#)

Tier 2: this quarter

5. **Host-level quality pass.** Resolve the healthcare PR split. Audit the programmatic index families page by page against one test: does this page say something the vendor's own site does not? Enrich or noindex accordingly. Internal linking: 20 to 45 contextual links with varied, descriptive anchors into each of the site's core commercial service pages, from Insights and sibling pages, not just navigation. Measure: indexed canonical sitemap URLs, impressions and clicks on the six pages, "Crawled, currently not indexed" count as a quality signal. [\[A\]](#) [\[C\]](#)
6. **Commercial pages that end the search.** Rewrite the core service pages against the Navboost target: scope, a price signal, process, proof, named partner, what happens in the first 48 hours. A buyer who finds the answer does not return to the results. Measure: Search Console CTR and, in PostHog, bounce-to-results and enquiry rate on those pages. [\[A\]](#)
7. **Adopt the article template in section 6** for every Insight from now, and retrofit the ten Insights with the most impressions. Claude-drafted with human bylines is within Google's rules provided the original inputs, verification and named accountability are real. Clinical and YMYL pieces keep the fortnightly review cap. Measure: per-URL AI impressions (Search Console), citation appearance rate in the prompt set, time to first fetch. [\[A\]](#) [\[C\]](#)

Tier 3: the year

8. One flagship recurring data asset. The strongest and most consistent evidence in the whole review is for recurring original research: 55 percent of decision-makers name research and data as the top quality marker, research reports are the most under-supplied format, 55 percent of journalists used PR-supplied research last year, annual surveys earn roughly ten times the links of comparable content. A firm in this position typically already has two or three candidate data assets sitting in production. Choose one, run it on a fixed cadence at a permanent URL that is updated rather than replaced, publish the method in full (sample, dates, sources, weighting, caveats), give it a one-sentence headline metric, byline a named specialist, and re-pitch prior linkers each edition. Expect two to three cycles before it is a citation. Measure: referring domains and mentions to the asset URL, journalist pickups, appearance rate in the prompt set. [c]

9. Named experts, deployed by evidence. An in-house technical or clinical specialist is the archetype Edelman finds 22 points more credible than a chief executive; byline relevant data releases under that name with a first-person method note and an interview offer. A senior figure with recognisable public-affairs experience sits low on public credibility scales but high on the "understands my specific challenges" factor public-affairs buyers weight at 85 percent; a short, regular, opinionated note on how the relevant political process will actually behave, with one contestable claim per edition, matches both the buyer evidence and the link behaviour of opinion pieces. Neither should be a generic brand ambassador. [c]

10. Earned mentions programme. Digital PR of the owned data assets only (never client assets), pitched under 200 words before noon with one follow-up. Real listings with three to five verified client reviews on Clutch, Sortlist, The Manifest and PRWeek, because those pages are the results page for head terms and are among the sources language models cite when asked who to hire; this needs owner-level account access. A YouTube presence built from existing podcast and interview footage is the one channel where the correlational evidence (0.71 to 0.74) justifies a low-cost trial. Never buy links. Measure: unlinked and linked mentions to the domain in Ahrefs, appearance rate in "who to hire" prompts. [c]

11. Distribution over production. Every Insight gets a personal LinkedIn post from its named author, a LinkedIn newsletter edition, an email to the list and a targeted pitch to the ten to twenty journalists on that beat. Personal profiles beat company pages in every dataset found; the multiples are soft, the direction is not. [c] [d]

What to stop or not start

- Extending llms.txt, building markdown mirrors, or adding FAQ schema. Any plumbing pilots already in scope should be closed rather than extended unless crawler logs show fetches, and the logs across 137,000 domains say they will not.
- Any content produced at scale without unique data per page, and any templated satellite page family beyond a small, already-capped set of near-duplicate regional pages.
- Single-run AI visibility scores as a KPI, and Domain Rating as a headline marker. The headline markers are followed referring domains and mentions to the six money pages and the flagship asset, segmented Search Console CTR on those pages, and appearance rate per engine with an interval.
- Reporting AI impressions as clicks, or AI referral conversion multiples as facts.
- Bumping dates, chasing word counts, chasing Core Web Vitals beyond the thresholds.

BACK MATTER

CONTESTED POINTS AND SOURCES

Unresolved

- Whether the fall in AI Overview citations from Google's top ten (76 to 38 percent) is mostly better parsing or mostly fan-out after the Gemini 3 upgrade. Ahrefs says both.
- What ChatGPT's web index is in 2026. OpenAI says third-party providers plus its own retrieval; an August 2026 researcher analysis suggests a proprietary index with 1.5 percent URL overlap against Bing. Unconfirmed.
- Whether brand mentions cause AI visibility or both follow from brand size. No interventional study exists.
- Whether schema influences AI Overview citation. Google says not required; Ahrefs' quasi-experiment says no; Microsoft says it helps its models understand content.
- Whether Google's claim that AI Overview clicks are "higher quality" holds. Google offers no data; Search Console gives site owners no AI click data to check it.
- The freshness contradiction between retrieval (fresher wins) and selection (older wins) in ChatGPT.
- Whether the worked-example site's click-rate gap on its highest-value page is explained by the snippet substitution. The experiment in section 10 answers it in four weeks.

Circulating claims that did not verify

- A June 2026 update to the Search Quality Rater Guidelines (live PDF is dated 11 September 2025; quoted headings do not exist).
- Edelman and LinkedIn 2025 reporting a buyer penalty for AI-sounding thought leadership (not in any accessible summary; PDF returned 403).
- "CEO posts get 7x the impressions of company pages" attributed to Edelman (appears to originate with Inc.).

- "62.1 percent of citations come from posts with question headings" and "40 to 61 percent of AI Overviews contain bullets" (vendor secondary claims with no locatable primary).
- HubSpot's "+106 percent from historical optimisation" (primary 404).
- Mueller quotes on schema for LLMs ("yes, no, it depends") and on digital PR being "as critical as technical SEO" (no dated primary post reached).
- The Q* 0.4 threshold for featured-snippet eligibility (not in the primary filing reached).

Primary sources consulted

Google and Microsoft

- AI features and your website (developers.google.com/search/docs/appearance/ai-features, updated 10 Dec 2025); Optimizing your website for generative AI features (developers.google.com/search/docs/fundamentals/ai-optimization-guide, updated 10 Jul 2026) and its announcement (Search Central blog, 15 May 2026); Top ways to ensure your content performs well in Google's AI experiences (21 May 2025).
- Creating helpful, reliable, people-first content; Using generative AI content; Core updates; Spam policies (updated 28 Aug 2026); March 2024 core update and spam policies post; Site reputation abuse (Nov 2024); Publication dates; Title links; SEO Starter Guide; Page experience; Crawl budget; Robots introduction; Block indexing; Consolidate duplicate URLs; Sitemaps; HTTP and network errors; Google common crawlers (updated 14 Jul 2026); JavaScript SEO basics; Structured data search gallery; Organization, Article, ProfilePage, FAQPage documentation; Multi-regional sites.
- Search Quality Rater Guidelines PDF, 11 Sep 2025 (static.googleusercontent.com/media/guidelines.raterhub.com/).

- Search Console Generative AI performance reports (Search Central blog, 3 Jun 2026; support.google.com/webmasters/answer/16984139); Search generative AI control (support.google.com/webmasters/answer/16908024; blog.google, 3 Jun 2026); Google Search Status Dashboard ranking history.
- blog.google: AI Mode update (20 May 2025); AI Mode in the UK (28 Jul 2025); Search at I/O 2026 (19 May 2026); Web Guide (24 Jul 2025); Google Marketing Live (20 May 2026); AI Overviews expansion (28 Oct 2024); Search On passage ranking (Oct 2020). How Search Works, ranking results.
- GA4 default channel groups (support.google.com/analytics/answer/9756891). Google patent US20200349181A1 / US11354342B2.
- Bing Webmaster blog: AI Performance public preview (10 Feb 2026); IndexNow adoption (Dec 2024); duplicate content and AI (19 Dec 2025); data-nosnippet (15 Oct 2025). Bing Webmaster Guidelines (26 Feb 2026, via Search Engine Roundtable and Search Engine Journal). indexnow.org FAQ. Fabrice Canel at SMX Munich (Search Engine Land, 20 Mar 2025).

Court record and leak

- US v. Google, Memorandum Opinion, Doc. 1033, 5 Aug 2024 (CourtListener), findings of fact 88, 92, 96, 97, 102, 103; remedies ruling 2 Sep 2025 (Knight-Georgetown Institute analysis; Hughes Hubbard; CNBC 5 Dec 2025).
- Content Warehouse API leak: SparkToro (28 May 2024); iPullRank (27 May 2024); Search Engine Land (28 May 2024, and Barnard on author entities, 6 Jun 2024).

Peer-reviewed and preprint

- Aggarwal et al., GEO, KDD 2024 (arXiv 2311.09735). Puerto et al., C-SEO Bench, NeurIPS 2025 (2506.11097). Kim et al., SAGEO Arena (2602.12187). Liu and Xu, FeatGEO (2604.19113). Martinez, critical survey of GEO 2023 to 2026 (2607.14035). Zhang, He and Yao, citation absorption (2604.25707). Xu, Iqbal and Montgomery, measuring AI Overviews (2605.14021). Allaham and Diakopoulos, Synthetic Sources (2605.23684). Yang, news citing patterns (2507.05301). Huang et al., Answer Bubbles (2603.16138). Kirsten et al., ACL 2026 Findings. Grossman et al., SIGIR 2026 (2604.27790). Schulte et al., Don't measure once (2604.07585). Zhang et al., source coverage and bias (2512.09483). Kumar and Palkhouski, GEO-16 (2509.10762). Nestaas et al., ICLR 2025. Cuconasu et al., EMNLP 2025 (2505.15561). Hutter et al., ECIR 2025. Liu et al., Lost in the Middle (2307.03172). Wan et al., ACL 2024.
- Pew Research Center, 22 Jul 2025. Tow Center / CJR, 6 Mar 2025. NIM, transparency without trust (2024 to 2025). Nielsen Norman Group, F-shaped pattern (reviewed 19 Aug 2026) and AI changing search behaviours (15 Aug 2025).

Industry studies

- Ahrefs: AI Overview citations vs rankings (Jul 2025; Mar 2026); brand correlations (26 May 2025; 12 Dec 2025); schema and AI citations (11 May 2026); llms.txt logs (15 Jun 2026); why ChatGPT cites pages (15 Apr 2026); fresh content (28 Jul 2025); AI Overviews reduce clicks (17 Apr 2025; Feb 2026 via PPC Land); AIO growth (13 May 2025); search traffic study (Dec 2023); backlink growth (2018); keyword difficulty (Dec 2025); Core Web Vitals (Jan 2025); topical authority (Jun 2026); featured snippets; meta descriptions; link-building case study; how long SEO takes.
- Semrush: AI Mode comparison (21 Jul 2025); most-cited domains (10 Nov 2025); AI Visibility Index (26 Jun 2026); ghost citations (9 Jun 2026); ChatGPT reasoning mode (30 Jun 2026); ChatGPT topic authority (20 Jul 2026). Growth Memo and Search Engine Land: how AI pays attention (16 to 18 Feb 2026); the consensus gap (11 May 2026); how AI picks its sources (Mar 2026). AirOps, The Fan-Out Effect (Apr 2026).

- Seer Interactive: AIO CTR updates (4 Nov 2025; 24 Apr 2026); content recency (25 Jun 2025). Amsive via Search Engine Land (21 Apr 2025). BrightEdge rank overlap (Sep 2025; Mar 2026). SE Ranking, Gemini 3 impact (26 Feb 2026). SparkToro zero-click (2 Jul 2024; 9 Jun 2026). Conductor AEO/GEO benchmarks (6 Jul 2026) and AIO volatility (14 Apr 2026). Similarweb gen-AI stats (29 Jul 2026). Adobe Q2 2026 via Search Engine Journal. Profound (Jun 2025). Peec (31 Mar 2026). Evertune via Search Engine Land (19 May 2026). Otterly (Feb 2026). Muck Rack Generative Pulse (7 May 2026). Writesonic (22 Jul 2026). Surfer (14 Jul 2026). GetMentions volatility (Jun 2026). Zyppy title rewrites (2022, updated 2026) and internal links (Feb 2026). McAlpin via Search Engine Land (1 May 2025). Backlinko ranking factors and CTR (Apr 2025). Sistrix CTR (2020, updated 2025). Portent (2020, 2021). Detailed.com (2025). HouseFresh (2024). Marsiglia Digital (May 2025). Lasso Security (24 Jun 2026). Studio36 UK AIO study (Oct 2025). Fractl AI statistics (Q2 2026).
- Vercel and MERJ, The rise of the AI crawler (17 Dec 2024). Vercel docs: bot management, attack mode, managed rulesets, BotID (2026). Cloudflare: Web Bot Auth (15 May 2025); Content Independence Day and pay per crawl (1 Jul 2025); Content Signals Policy (Sep 2025); crawlers and referrals (29 Aug 2025); AI Crawl Control docs; TechCrunch on the September 2026 default (1 Jul 2026). OpenAI bots documentation; Anthropic crawler documentation; Perplexity bots documentation; Apple, Amazon, Meta, Common Crawl crawler pages. IImstxt.org. Seekio and SEO Depths IImstxt log studies (2026). IETF AIPREF working group and draft-ietf-aipref-vocab-07. EU GPAI Code of Practice summaries. Next.js JSON-LD and sitemap documentation.

B2B and PR research

- Edelman and LinkedIn B2B Thought Leadership Impact Report 2024 and 2025 (landing pages; LinkedIn write-up; Curzon, Roo & Eve, ContentGrip and Demand Gen Report summaries). Edelman Trust Barometer 2024 and 2026 (summaries; PDFs returned 403). Hinge Research Institute: High Growth Study 2025 and 2026; Inside the Buyer's Brain; Referral Marketing (2015); Visible Expert. CMI/MarketingProfs B2B Content Marketing benchmarks 2025 and 2026. Cision State of the Media 2025 and 2026. Muck Rack State of Journalism 2025 and 2026. Orbit Media blogger survey 2025 and ten-year retrospective. Backlinko and BuzzSumo 912 million posts (2019). BuzzSumo and Majestic (2016). BuzzSumo and Mantis original research survey (2020). Fractl campaign studies and survey design guide (30 Apr 2026). Animalz EVE framework, copycat content, content refresh, benchmark report. LinkedIn B2B Institute 95-5 rule. 6sense Buyer Experience Report 2024 and 2025. Dreamdata journey benchmarks (2025). Refine Labs hybrid attribution. Gartner 2026 buyer surveys via Demand Gen Report.

Worked-example site baseline

- Ahrefs API, 2 Sep 2026: gsc-keywords and gsc-pages, 1 Jun to 31 Aug 2026; site-explorer-ai-responses-count and batch-analysis (GB) for the example site and three anonymised competitors. Bright Data Google GB results for "b2b pr agency london", "crisis communications agency", "healthcare pr agency". Live fetches of 15 site pages as Googlebot. The site's own front-end source code, inspected directly, with its git history.

Published by Bridgehead Communications, 2 September 2026. The per-claim research dossiers behind each grade, with URLs, dates and sample sizes, are available on request. Where this report and a dossier differ in grading, the more conservative grade applies; in particular the Limy blog post is graded as a commercially motivated summary of the Martinez survey, not as peer-reviewed work in its own right.

TALK TO US

We run the programme in section 10 on our own site, in public, and report the result the same way we'd report a client's. If you want your organisation's organic and AI-answer visibility read this rigorously, talk to our AI visibility practice (bridgeheadcommunications.com/ai-visibility).