

AEO FOUNDATIONS

Answer Engine Optimization: being the passage that featured snippets, People Also Ask, AI Overviews, and voice assistants extract and attribute. Ranking is no longer the prize; being the answer is.

IN THIS GUIDE

- 01 Why answers beat rankings

- 02 SEO, AEO, GEO: one foundation, three layers

- 03 How answer selection actually works

- 04 Drawing the lines: the taxonomy

- 05 Question-first content architecture

- 06 Schema and the technical layer

- 07 The llms.txt debate, honestly

- 08 Measuring answer share

- 09 The playbook: do this, in order

- 10 The mistakes that cost the most

- 11 The head game: faster signal, fuzzier payoff

- 12 What nobody knows yet · Sources

WHY ANSWERS BEAT RANKINGS

IN PLAIN WORDS

Search has changed. More and more people get their answer right there on the results page or from an AI, so they never click through to a website. That means being the short answer that gets shown, and named as the source, now matters more than just sitting high on the list of links.

The economics of search visibility inverted between 2024 and 2026. SparkToro's like-for-like clickstream series puts zero-click searches at 60.45% of US Google searches in 2024 and 68.01% by early 2026. Pew Research's behavioral study of 900 US adults found that when an AI summary appears, users click a traditional result on only 8% of visits versus 15% without one, and click a link inside the summary just 1% of the time. Ahrefs measured the CTR penalty for the top result at 58% when an AI Overview is present, worse than the 34.5% it measured in April 2025.

8% vs 15%

Visits where users click a traditional result with vs without an AI summary present (Pew Research, 68,879 searches, 2025).

37.9%

Share of AI Overview citations from top-10 organic results in early 2026, down from 76% in July 2025 (Ahrefs, 863K SERPs).

40-60

Words in the answer-first block that snippet and answer systems most reliably lift (Moz snippet analysis; set-wide standard).

The other side of the ledger, from the SEO guide's canonical narrative: traffic is repricing, not disappearing. Seer Interactive's 5.47-million-query dataset shows organic CTR on AI Overview queries rebounding 85% off its December 2025 floor, and pages cited inside the answer earning roughly 120% more clicks per impression than uncited pages (with the caveat that the rebound may reflect a temporary UI change). Citation is the new position one, and AEO is the craft of winning it.

ONE FOUNDATION, THREE LAYERS

IN PLAIN WORDS

There are three related jobs that stack on top of each other. The first gets your page found and ranked, this middle one gets one small piece of your page picked as the answer, and the third gets your whole brand known and recommended by AI tools. You cannot skip the lower steps and still win the higher ones.

SEO

THE RETRIEVAL LAYER

UNIT: THE PAGE · rank & get clicked

AEO

EXTRACTION FROM THE INDEX

UNIT: THE PASSAGE · be the cited answer

GEO

RECOMMENDATION ACROSS ENGINES

UNIT: THE BRAND · be named & recommended

ChatGPT · Perplexity · Claude · Gemini · AI Overviews

AEO is the middle layer: you have SEO's retrieval, now you compete to be the passage the answer surface extracts.

SEO, AEO, and GEO are concentric layers on one shared foundation, not rival disciplines. **SEO** is the retrieval layer everything else sits on: every major answer and generative engine retrieves from a conventional crawler-built index. Google's index feeds AI Overviews, AI Mode, and Gemini; Bing's feeds ChatGPT search and Copilot; Perplexity runs its own crawler. **AEO** sits on top and changes the unit of optimization from the page to the passage: being the specific chunk an answer surface extracts and attributes. **GEO** sits on top of both and changes the unit again, from the passage to the brand: what generative engines already know about you, plus earned presence in the third-party sources they retrieve and cite.

A page cannot win AEO while failing SEO basics, and GEO without AEO-shaped content gets retrieved but not cited. AI Overviews are the surface where all three meet.

READ THIS GUIDE IF

You rank but are not the extracted answer. If you first need to be retrievable and ranked, read **SEO Foundations**. If you need LLMs to name and recommend your brand, read **GEO Foundations**. All three at selwynuy.dev/guides.

HOW ANSWER SELECTION WORKS

IN PLAIN WORDS

This explains where those top answers actually come from. Search engines take the boxed answer at the top, called a snippet, from pages that already rank near the top, and AI tools quietly break your one question into several smaller ones and grab a short passage to answer each. Knowing this shows you that you have to rank first, and that every little chunk of your page can get picked on its own.

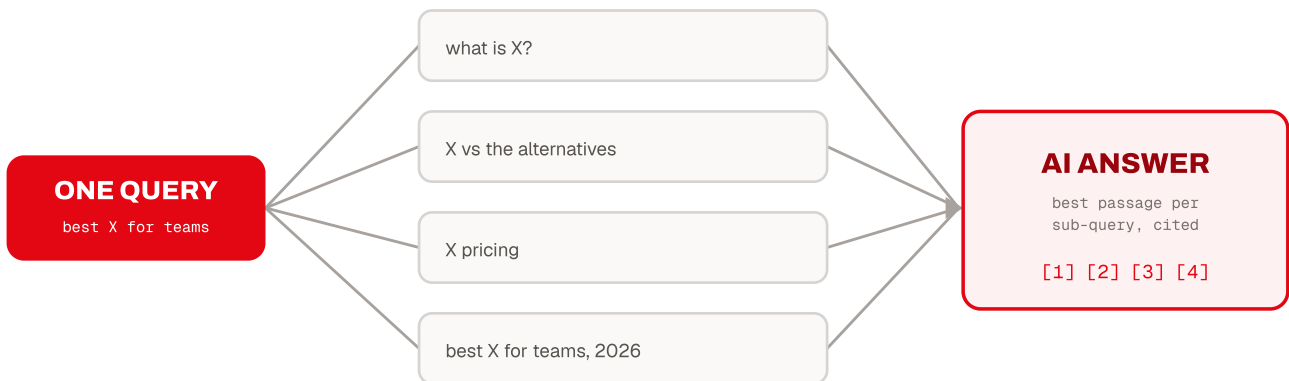
Featured snippets are an elevation, not a separate index

Google's systems decide whether a page "would make a good featured snippet" and elevate it from the normal ranking set. Ahrefs found 99.58% of snippet winners already rank top-10 for the query, so you cannot snippet-optimize your way around weak rankings. Snippets are won by page-one pages that present the most extractable, information-dense passage; winning paragraph answers cluster around 40 to 50 words (Moz, 1.4M snippets).

People Also Ask cascades

PAA is a graph of related questions, each carrying its own snippet. Winning one slot frequently cascades: the same passage serves the PAA accordion, the standalone snippet for that question, and voice answers.

AI Overviews and AI Mode retrieve passages



Query fan-out: one query becomes many, a passage is retrieved for each, and the answer cites the best. You get cited for questions the user never typed.

Google's generative surfaces use query fan-out: one query becomes many synthetic sub-queries, candidate passages are retrieved for each, and a Gemini model synthesizes a grounded, cited answer (the GEO guide owns the full pipeline mechanics). Two consequences: citation happens at the passage level, so a page is a portfolio of independently liftable chunks; and the retrieval net is widening beyond classic rankings, with top-10 results supplying 37.9% of citations in early 2026, down from 76% in mid-2025.

Voice is won as a byproduct

Google Assistant reads featured snippets aloud (Backlinko measured 40.7% of voice answers sourced from snippets); Alexa leans on Bing plus Wikipedia and Wolfram Alpha. Clean, speakable, short-sentence answers win voice without separate work.

GOOGLE'S POSITION

"There are no additional requirements to appear in AI Overviews or AI Mode" beyond being indexed and snippet-eligible. True but incomplete: eligibility is table stakes. Selection is the competition.

04 - THE TAXONOMY

DRAWING THE LINES

IN PLAIN WORDS

This is a clear table that lays those three jobs side by side, what each one is trying to do and how you measure it. People throw these terms around loosely and mix them up, so having one plain definition to point to keeps you from getting confused or oversold.

The industry blurs these terms constantly, so this table is the set-wide reference (the SEO and GEO guides point here instead of drawing rival taxonomies).

	SEO	AEO	GEO
Objective	Rank and get clicked	Be extracted and attributed	Be known and recommended
Surface	Ranked results, rich results	Snippets, PAA, AI Overviews / AI Mode citations, voice	Generative answers across ChatGPT, Perplexity, Claude, Gemini, AI Overviews
Unit	The page	The passage	The brand / entity (plus the passage)
Primary levers	Crawlability, relevance, links, page experience	Question-first structure, answer-first blocks, format matching	Earned mentions, entity consistency, training-data presence, citability
Measurement	Positions, organic sessions	Answer share per surface	Visibility rate across sampled runs
Time horizon	Months	Weeks to months	Months to model-release cycles

These are concentric circles, not rivals. SEO supplies retrieval; AEO shapes the passages retrieval systems extract; GEO adds the off-site reputation layer generative engines consult. AI Overviews are explicitly the overlap surface where AEO's passage work and GEO's brand work meet on top of SEO's index work. One warning the whole set repeats: vendors use AEO, GEO, LLMO, and AIO interchangeably. Define terms once and hold the line.

QUESTION-FIRST ARCHITECTURE

IN PLAIN WORDS

This is the hands-on writing part. Find the real questions people ask, put each question as a heading, then right under it give a short answer of about forty to sixty words that makes full sense on its own, and only after that add the longer details. Write each chunk so it still reads clearly if someone lifts it out by itself, because that is exactly what a search engine or AI does.

Question research replaces keyword research

Mine People Also Ask graphs with AlsoAsked, use keyword tools' question filters, and pull from your own site search and support tickets. Then sample the real answers: ask ChatGPT, Perplexity, and AI Mode the questions your buyers ask, and record who gets cited and why. Cluster questions by intent stage; map each cluster to one URL.

The answer-first block, the atomic unit of AEO

How much does X cost?

H2 · the literal question

X starts at \$19 per user per month on the Team plan, billed annually.
A free tier covers up to three seats. Enterprise pricing is custom.

40-60 words · stands alone · restate the subject, no pronouns

then: evidence, nuance, examples, edge cases, for the human who clicks through

The atomic unit of AEO: a question heading, a 40-60 word answer that stands alone, then the depth. The answer block is what gets lifted into the snippet.

- An H2 or H3 phrased as the literal question a user asks ("How much does X cost?", not "Pricing considerations").
- Immediately below: a direct, self-contained answer of roughly 40 to 60 words, never more than 80, that makes sense with zero surrounding context.
- Then expand: evidence, nuance, examples, and edge cases for the human who clicks through. The inverted pyramid, applied per section rather than per page.

Match format to snippet type

Paragraphs for definitions and "what is" questions; real `` markup for how-to steps; tables for comparisons, pricing, and specs; short declarative sentences for voice.

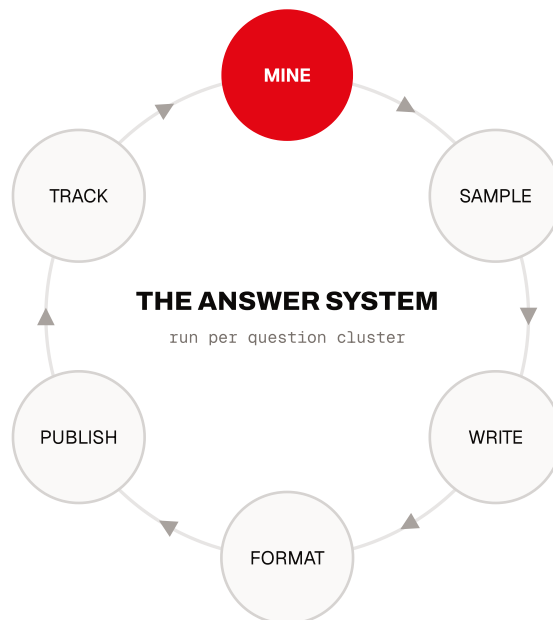
Passage independence

Because fan-out retrieves chunks against machine-written sub-queries, every section should survive being read in isolation: restate the subject noun instead of leaning on pronouns, keep one idea per section, and cover the adjacent sub-questions (cost, alternatives, timeline, risks) on the same page or a tightly linked cluster. Give each question heading a stable anchor link so engines can deep-link the passage.

GOOGLE'S COUNTERPOINT, QUOTED

Google's May 2026 guide says publishers "can ignore tactics like chunking content, creating unnecessary AI text files, or pursuing inauthentic mentions" because its systems understand multi-topic pages. Keep the passage-first writing anyway, for two honest reasons: it is simply clear writing, and non-Google pipelines (ChatGPT, Perplexity) still retrieve chunks. Frame it as craft, not as a Google hack.

The answer system



A repeatable cycle per question cluster: mine the real questions, sample who wins today, write the answer block, match the format the SERP shows, publish, then track citation share and feed it back.

The system is small enough to hand off. Mining and sampling produce a question list with a "who owns it now" note; writing fills the answer-block template; formatting matches the live surface; tracking waits for repeated samples before judging. The work is never a one-off audit because the answer surface shifts under you.

The answer-block checklist

The QA a writer or VA runs on every answer block before it ships.

Heading is the literal question a user asks	Matches the query and is what the engine extracts
Answer is 40-60 words, self-contained	Fits the snippet and survives being lifted out
Subject restated, no pronouns	The passage works when read in isolation
Format matches the live SERP (para, list, table)	The engine picks the format it already shows
A stat, quote, or cited source in the passage	The KDD-validated citation boosts
A stable anchor link on the heading	Engines can deep-link the passage
Organization / Person schema present	Entity clarity for citation eligibility
Page is server-rendered (JS-off check)	Most AI fetchers execute little or no JavaScript

06 - THE TECHNICAL LAYER

SCHEMA AND ACCESS

IN PLAIN WORDS

Schema is a set of hidden labels you add to a page that tell search engines what things are, like this is the company, this is the author. These labels will not magically get you picked, but they help engines trust who you are. The bigger thing here is plumbing: make sure the AI programs that read the web are actually allowed onto your site, because they cannot quote a page they are blocked from reading.

What Google actually says

No special markup, no AI text files, and no special structured data are required for AI Overviews or AI Mode. Schema is not a citation cheat code. Its real value is entity disambiguation and the remaining rich results.

The FAQ and HowTo reality check

The canonical timeline: HowTo rich results ended September 2023; FAQ rich results were restricted to government and health sites in August 2023, stopped appearing in May 2026, lost search-appearance and Rich Results Test support in June 2026, and lose Search Console API support in August 2026. FAQPage markup remains valid schema.org and causes no penalty; it just buys nothing visually. Q&A-formatted *content* remains one of the most extractable patterns. Format beats markup.

Schema that still earns its keep

Organization and Person with sameAs links to Wikipedia, Wikidata, and LinkedIn (these feed the knowledge graphs answer engines consult); Article with author attribution; Product, Offer, and Review for shopping answers; LocalBusiness for voice and local; Speakable (still beta, expectations low); QAPage for genuine forum content.

Crawler access is the real battleground

Answer engines cannot cite what they cannot fetch. The GEO guide carries the canonical crawler taxonomy table; the two facts to carry here: Google-Extended controls Gemini training, not AI Overviews (which use normal Googlebot), and OpenAI's training bot (GPTBot) is a different token from its search bots (OAI-SearchBot, ChatGPT-User). Many sites blocked AI bots at the CDN in 2023-2024 and are still silently opted out of citations. Also: nosnippet, data-nosnippet, and max-snippet limit what Google can quote in snippets and AI Overviews alike. You can protect content from summarization or compete for citation, not both.

Rendering: most AI fetchers execute little or no JavaScript. Server-render the answer blocks. This is the SEO guide's territory; its rendering, soft-404, and canonicalization rules are prerequisites here.

07 - THE PLACEBO

THE LLMS.TXT DEBATE

IN PLAIN WORDS

There is a special file called llms.txt that some people push as the secret to getting picked by AI. The honest truth is that almost no AI actually reads it, so it does next to nothing except for a few sites that are mostly documentation. The real lesson is to trust what search engines actually do, not what salespeople wish they did.

llms.txt is the most hyped and least effective artifact in AEO, and this guide gains credibility by covering it honestly. The proposal (Jeremy Howard, September 2024): a markdown file at /llms.txt listing a site's most important pages in LLM-friendly form. Pitched as robots.txt for the AI era, inverted into an invitation.

The case against, with receipts

- Google does not use it. John Mueller, June 2025: "no AI system currently uses llms.txt", with his comparison to the long-dead keywords meta tag: self-declared, unverified, trivially gameable. Gary Illyes confirmed no support and no plans.
- Ahrefs analyzed ~137,000 domains; of the ~38,000 with a valid llms.txt, 97% received zero requests for the file in the measured month. Adoption is real; consumption mostly is not.
- No major answer engine has committed to it as a ranking or citation input.

The case for, narrower than advocates claim

- Some AI coding tools fetch it opportunistically; documentation-heavy sites report retrieval by coding assistants where a curated markdown index genuinely helps an agent navigate docs.
- Cost is near zero if generated automatically.

VERDICT

A cheap, low-priority hedge for documentation-heavy sites; irrelevant for everyone else. It is not AEO. Anyone selling llms.txt as the centerpiece of AI visibility is selling a placebo. The transferable lesson: always separate what engines document and demonstrably do from what the industry wishes they did.

08 - SCOREBOARD

MEASURING ANSWER SHARE

IN PLAIN WORDS

Answer share is a simple scorecard: out of the questions that matter to your business, how often are you the answer that gets shown or credited. Because AI gives a slightly different answer each time you ask, you have to check the same questions many times over weeks and watch the trend, not trust a single day's reading.

Define **answer share**: of a fixed basket of questions that matter to the business, the percentage where you are the displayed answer or a cited source, per surface (snippet, PAA, AI Overview, AI Mode, ChatGPT, Perplexity), tracked over time against named competitors.

First-party signals

- Search Console folds AI Overview and AI Mode activity into "Web" data with no breakout. State the limitation plainly: you cannot isolate AI Overview traffic in GSC.
- Referral segments for chatgpt.com, perplexity.ai, copilot.microsoft.com, gemini.google.com. Volumes are small; measure your own conversion rates before repeating vendor claims.
- Server logs: AI crawler hits (GPTBot, OAI-SearchBot, PerplexityBot, ClaudeBot) are a leading indicator of retrievability.

Tools

Rank trackers (Semrush, Ahrefs, seoClarity) report snippet and PAA ownership per keyword, still the cheapest reliable answer-share data. AI visibility platforms (Profound, Peec AI, Otterly.AI, Semrush AI Toolkit, Ahrefs Brand Radar) all work the same way: run a prompt basket against multiple LLMs on a schedule, parse who gets mentioned and cited, compute share of voice. The full tool stack, with cost tiers and a setup order, lives in The Stack companion.

METHODOLOGICAL HONESTY

LLM answers are non-deterministic: the GEO guide's canonical evidence (SparkToro and Gumshoe, 2,961 runs) found under a 1-in-100 chance of the same brand list twice, yet stable 55-77% inclusion rates for top brands. Insist on repeated sampling and trend lines; treat every single-day reading, and every "AI rank #3" dashboard, as noise.

DO THIS, IN ORDER

IN PLAIN WORDS

This is the whole thing turned into a checklist you can follow from the top. It is split into three levels, from just getting allowed and eligible, to rewriting your pages to win the answer, to measuring it and keeping it fresh. Each step tells you exactly what to do and how to know it worked, so a beginner can just start at step one.

The whole discipline as an ordered plan. Follow it top to bottom if you are starting out; jump to your level if you are not. Every step names the tool and how you know it worked.

LEVEL 1 STARTER

become eligible

- 1 Confirm you are snippet-eligible.**

Search your codebase and meta tags for accidental `nosnippet`, `data-nosnippet`, or a low `max-snippet` on pages you want quoted.

Done when: the pages you care about carry no snippet-suppressing directives.
- 2 Check which AI crawlers you allow.**

Open `robots.txt` and your CDN/WAF rules (Cloudflare). Confirm OAI-SearchBot, PerplexityBot, and Bingbot are not blocked.

Done when: a fetch as each agent (or a log check) returns 200, not 403.
- 3 Build your first question list.**

Mine `a1soasked.com` or your keyword tool's Questions filter, and add the questions from your own support tickets.

Done when: you have 20-plus real questions, written as questions.
- 4 Sample who wins today.**

Ask ChatGPT, Perplexity, and Google AI Mode your top questions and record who gets cited.

Done when: you know who currently owns each answer.

- 5 Retrofit your page-one pages first.**
- For each page already ranking top-10, add a question-phrased H2 and a self-contained 40-60 word answer beneath it.
- Done when:** the answer block reads correctly with zero surrounding context.
- 6 Match format to the answer type.**
- A paragraph for definitions, a real for how-to steps, a table for comparisons.
- Done when:** your format matches what the live SERP shows for that query.
- 7 Give every question heading an anchor link.**
- Add id anchors so engines can deep-link the passage.
- Done when:** each heading is individually linkable (yoursite/page#the-question).
- 8 Add entity schema.**
- Render Organization and Person JSON-LD with sameAs links to your real profiles.
- Done when:** the Rich Results Test validates it.

- 9 Define an answer-share basket.**
- Fix 50-100 questions that matter and, per surface (snippet, AI Overview, ChatGPT, Perplexity), record whether you are the answer.
- Done when:** you have a baseline percentage per surface.
- 10 Track it with repeated sampling.**
- Use an AI-visibility tool (Profound, Peec AI, Otterly) or your own prompt script, running each prompt many times per week.
- Done when:** you have a trend line, not a single-day reading.
- 11 Move off-site where answers are cited.**
- Earn presence on the surfaces AI answers pull from: Reddit, YouTube, review sites, credible publications.
- Done when:** your brand appears in the cited sources, not only on your own domain.
- 12 Refresh on a cadence.**
- Re-sample the basket monthly and update answer blocks when facts age.
- Done when:** stale facts (pricing, dates) never survive a monthly pass.

Effort allocation, honestly: roughly 60% content restructuring, 20% measurement, 15% off-site, 5% technical files and markup.

THE MISTAKES THAT COST THE MOST

IN PLAIN WORDS

This is the list of common mistakes that burn the most time. The big ones are trying to win the answer on pages that do not even rank yet, hiding your answer under waffle, and trusting a single day's result or shiny vendor claims. Reading this once saves you from the traps most people fall into.


- Chasing answer placement on pages that rank nowhere. Snippets elevate from page one; AEO on unranked pages is building on sand.
- Treating llms.txt, FAQ schema, or any machine-readable file as the strategy. Format and substance beat markup, and the files are unread.
- Burying the answer under preamble, hedges ("It depends on many factors..."), or marketing throat-clearing instead of a self-contained 40-60 word block.
- Writing pages instead of passages: pronoun-dependent, interdependent sections are unextractable even on excellent pages.
- Blocking AI crawlers indiscriminately, or confusing Google-Extended with Googlebot and GPTBot with OAI-SearchBot.
- Reporting single-day AI visibility readings as fact.
- Repeating unverifiable vendor statistics ("FAQ schema lifts snippets 35%") without tracing them to a real study. This is the discipline's biggest credibility trap.
- Measuring success in clicks alone when 68% of searches end without one. Track citation share, mention sentiment, and answer share.

FASTER SIGNAL, FUZZIER PAYOFF

IN PLAIN WORDS

This is about what to expect and how to stay calm. You will see results faster than with regular search, sometimes in a week or two, but there is no steady rank to watch, so a mention can show up and then vanish. Judge success by whether the AI quotes you accurately and often over time, not by a single good or bad day.

AEO rewards you sooner than SEO, but on a surface with no stable rank to grip. The mental shift is from "what position am I" to "am I the answer, and did it get my facts right".



AEO moves faster than SEO: first citations in one to two weeks on a strong domain, but consistent citation and real business impact still take three to six months. The payoff is presence, not clicks.

What to expect

Technical changes (structure, schema, answer-first blocks) can show in two to four weeks on an established domain; first citations in one to two weeks for strong sites, four to six for new brands; consistent citation patterns in three to six months. Success looks like the machine quoting you accurately and often, not a traffic spike.

Traps to avoid

1

Importing the rank-check reflex

An AI citation is not a keyword position; it appears and disappears. Checking it like a rank produces false alarms.

2

Over-reading a first-week win

Early citations arrive fast, so teams extrapolate one good week into a trend, then panic when it vanishes. Sample over weeks.

3

Formatting tricks over trust

Citation eligibility runs on E-E-A-T and verifiable accuracy. Chasing snippet hacks while ignoring authority underperforms.

THE METRIC THAT MATTERS

Judge AEO on citation consistency and the factual accuracy of the quote, not click volume. Treat it as a recurring workflow, not a one-off audit (Aleyda Solis): the answer surface shifts under you, so the work is never finished.

WHAT NOBODY KNOWS YET

IN PLAIN WORDS

This is an honest admission of what is still unknown. Nobody fully knows how the AI picks which sources to credit, the numbers get argued over, and the tricks change fast. So stay humble, note the date on anything you read, and put your effort into the things that last: genuinely helpful content, being trusted, staying readable to crawlers, and keeping a habit of measuring.

- **Selection mechanics are partially opaque.** Google documents fan-out but not how citations are chosen among grounded candidates. The top-10 overlap falling from 76% to 37.9% in under a year shows the ground shifts fast enough to invalidate playbooks mid-year. Date every claim.
- **The zero-click math is contested.** Clickstream panels differ; Google disputes the framing without publishing data. Present ranges.
- **Does schema help AI citation?** Google says no special structured data is needed. Vendor lift numbers rarely share methodology. The defensible claim is entity disambiguation, not bought citations.
- **Non-determinism undermines measurement.** No tool observes real user prompt distributions; all sample. Whether tracked baskets represent actual buyer behavior is unsolved.
- **Attribution of business value is immature.** "AI traffic converts 4-6x better" circulates widely and traces to vendor posts, not audited data.
- **Durability risk.** Google restricted FAQ rich results overnight in 2023 and reshaped citation patterns within months in 2025-2026. AEO tactics have shorter half-lives than SEO tactics. The durable assets: content that genuinely answers questions, entity authority, crawl access, and a measurement habit. The GEO guide's contested-tactics list continues this thread.

Sources

- Google Search Central: AI features and your website developers.google.com/search/docs/appearance/ai-features
- Google Search Central: featured snippets developers.google.com/search/docs/appearance/featured-snippets
- Google Search Central: optimizing for generative AI features (May 2026) developers.google.com/search/docs/fundamentals/ai-optimization-guide
- Pew Research Center: clicks with AI summaries (2025) pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/
- Ahrefs: AI Overviews reduce clicks by 58% ahrefs.com/blog/ai-overviews-reduce-clicks-update/
- Ahrefs: 38% of AI Overview citations pull from the top 10 ahrefs.com/blog/ai-overview-citations-top-10/
- Ahrefs: 97% of llms.txt files never get read ahrefs.com/blog/llmstxt-study/
- SparkToro: 2026 zero-click study sparktoro.com/blog/in-2026-less-than-one-third-of-google-searches-still-send-a-click/
- Aggarwal et al.: GEO, ACM KDD 2024 arxiv.org/pdf/2311.09735
- Search Engine Journal: Google drops FAQ rich results (2026) searchenginejournal.com/google-drops-faq-rich-results-from-search/574429/
- Search Engine Journal: llms.txt has no current implementation searchenginejournal.com/google-says-llms-txt-is-purely-speculative-for-now/577576/
- Search Engine Land: query fan-out guide searchengineland.com/guide/query-fan-out
- Seer Interactive: AIO impact on CTR, 2026 update seerinteractive.com/insights/aio-impact-on-google-ctr-2026-update
- Backlinko: voice search study backlinko.com/voice-search-seo-study
- Digiday: WTF are GEO and AEO digiday.com/media/wtf-are-geo-and-aeo-and-how-they-differ-from-seo/
- Zapier: AI visibility tools, 2026 landscape zapier.com/blog/best-ai-visibility-tool/