



ANSWER ENGINE OPTIMISATION

The AEO Playbook.

AI tools now answer questions instead of listing websites. This guide explains how to get your business included in those answers.

A briefing for the people
who sign off the budget

14 pages

15 min read

What’s inside

01	Why AEO matters now	03
02	Where the answers come from	04
03	How answer engines read you	05
04	One question, twenty searches	06
05	Seven things to do	07
06	Measuring what actually matters	08
07	Four myths worth retiring	09
.	Benchmark: 712 B2B companies audited	10
08	Quick wins for the next 30 days	11
.	Worksheet: the AI-readiness audit	12
.	Sources and method	13

SEO gets you ranked.

AEO gets you quoted.

You still need both. This book is about the second one. The SEO Playbook covers the first.



Why AEO matters now

Your customers have changed how they search. They ask a full question and get a written answer, put together from several websites, often without clicking any of them.

So the question is no longer where you rank. It is whether you are in the answer. If a competitor is quoted and you are not, the buyer may never know you exist.

THE SHORT VERSION

Traffic is no longer the scoreboard

Expect fewer sessions, higher intent per session, and influence you can only see by asking the AI tools what they say about you.

Four numbers worth knowing

INDEPENDENT STUDIES, 2025 TO 2026

68%

of US Google searches ended without any click in early 2026, up from about 45% a decade ago.¹

8%

click rate when an AI summary is shown, against 15% without one, across 68,879 tracked real searches.³

4.4x

the conversion rate of AI-referred visitors compared with traditional organic visitors.⁶

+35%

organic click-through for brands cited inside an AI Overview, against uncited brands on the same page.⁵

IN SHORT Fewer visits, but better ones, and a clear benefit to being named. Every figure in this book is numbered to its source on page 13.



Where the answers come from

Five assistants shape most AI-mediated discovery. They differ in where they get their facts, and that decides how you earn a mention. Two of them read only static HTML.¹⁰

ChatGPT

Conversational · Bing-fed

Over a billion weekly users, and the least traceable citation surface: it often paraphrases without naming anyone.⁹ **Lever:** rank in Bing, write in clean answer blocks, state your expertise on the page.

Gemini & AI Mode

Google index · inside Search

Gemini passed a billion monthly users, and the same model powers AI Overviews across 2.5 billion monthly users of Search.⁹ It is also the only major engine that reliably renders JavaScript.¹⁰ **Lever:** schema, author credentials, fast mobile pages.

Claude

Enterprise-heavy · static HTML

Small in consumer share but the fastest riser, and heavily used inside businesses for research and drafting.⁸ Reads static HTML only. **Lever:** server-rendered pages, documentation-grade clarity, credible third-party sources.

Perplexity

Citation-first · 5 to 12 sources

Every answer carries visible links, so a citation is genuine exposure. It shares only about 11% of its top-cited domains with ChatGPT: the two read different webs.¹¹ **Lever:** hard facts, comparisons, roundups and forums.


Microsoft Copilot

Bing-fed · inside Windows & Office

Modest on the open web but embedded where office work happens. **Lever:** none of its own. The Bing work you do for ChatGPT pays here twice.

Share of assistant web visits

WORLDWIDE, MAY 2026⁸



Bars are relative to ChatGPT, not to all AI use. Apps add another 15 to 60% on top depending on the assistant, and AI Overviews sit outside this chart entirely, inside a product handling billions of searches a day.⁸

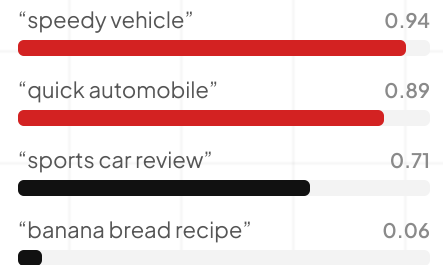
How answer engines read you



These systems do not match your words. They convert language into numbers, called embeddings, that place ideas near each other by meaning. “Fast car” and “speedy vehicle” land in almost the same spot. Repeating a keyword adds nothing. Being unmistakably clear adds everything.

When a model needs current facts it retrieves the passages closest in meaning to the question, then writes from them. Your page is not judged as a page. Single paragraphs compete on their own.

FIG. 1 · CLOSENESS IN MEANING TO “FAST CAR”



Illustrative values, not measured data.

01 RETRIEVE

The engine searches meaning, not text, and grabs the nearest passages.

02 SELECT

Each candidate is scored for clarity, trust and fit with the question.

03 ASSEMBLE

The best passages are combined into one answer, with or without credit.

What this rewards

- / Open with a 40 to 60 word summary that states the answer.
- / Name real things: products, places, standards, people.
- / Show authorship, dates and credentials.
- / Define your terms plainly. “X is...”
- / Use lists, tables and Q&A so passages lift cleanly.
- / Link related pages so the topic reads as a network.

TIP If a paragraph only makes sense after reading the three above it, it is unlikely to be quoted. Write each section so it stands on its own.



One question, twenty searches

A modern answer engine rarely searches the phrase it was given. It splits the request into sub-questions, runs them in parallel, and merges the best passage for each. Like a journalist quoting five experts to file one story.

WHAT WAS ASKED

“Best laptops for students under \$1,000 with long battery and backlit keys”

WHAT WAS ACTUALLY SEARCHED

- / Laptops under \$1,000 for study
- / Battery life compared, student models
- / Which laptops have backlit keyboards
- / Anything meeting all three at once

Answer one part well

You do not need to win the whole query any more. You need to be the best answer to one part of it. Stop writing the single page that chases the full phrase. Publish a small cluster where each piece owns a facet, tied together by a hub.

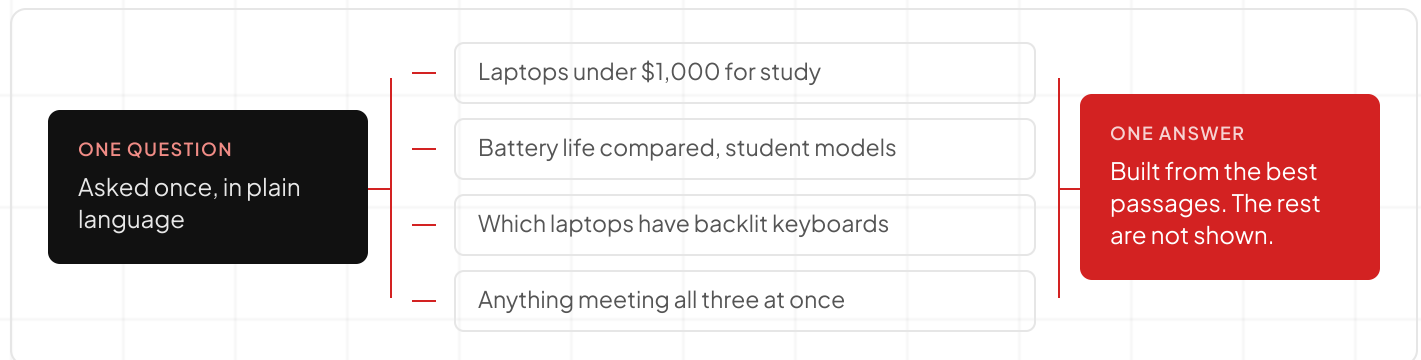


FIG. 2 · QUERY FAN-OUT, NAMED BY GOOGLE IN ITS OWN SEARCH BLOG¹⁵

TIP List the ten questions a buyer asks before choosing you. Every one your site does not answer directly is a question someone else is being quoted for.



Seven things to do

- 01 Make the site agent-friendly**
Agents now book, buy and fill forms for users. Run a real task through one. If an agent cannot finish checkout, that revenue is unreachable.
- 02 Ship complete structured data**
Schema states what a page is. Audit for the types you are missing: product, service, location, event, course, FAQ.
- 03 Answer questions, not keywords**
Target the long, specific phrasing people actually use, and answer the situation behind it instead of pitching.
- 04 Widen the footprint**
Answers get assembled from forums, video, community threads and press. More credible places means more chances to be pulled in.
- 05 Fix the technical floor**
Fast, mobile-clean, crawlable, readable without JavaScript, not blocking the crawlers you want. Unreachable content cannot be cited.
- 06 Keep the indexing basics tight**
Verify in Search Console and Bing Webmaster Tools, submit sitemaps, push new pages for indexing. AI answers still stand on old plumbing.
- 07 Build what a summary cannot replace**
Calculators, selectors, checklists, your own data. An answer box can restate your blog post. It cannot run your pricing tool.

If an AI agent cannot complete a purchase on your site,
neither can the customer it is acting for.

Measuring what actually matters



Falling traffic is not always a bad sign. Much of it was never going to buy. The more useful question is whether you appear in the answer when someone is deciding. Publisher pageviews from Google fell 38% across US sites in a single year⁷, while AI Mode refers on just 1.6 to 2.5% of queries against 17 to 19% for classic results.²

CLICK RATE, WITH AND WITHOUT AN AI SUMMARY³

15%



No AI
summary

8%



AI summary
shown

Across 68,879 tracked searches. Only 1% of visits clicked a link inside the summary, and 26% of sessions ended there.³

Five metrics to adopt

- 01 **AI citations.** Are you named when the tools answer your category's questions?
- 02 **Task completion.** Can an agent finish a real job on your site?
- 03 **Cross-platform presence.** Do you appear where the engines look?
- 04 **Schema accuracy.** Is your structured data valid and complete?
- 05 **Agent-assisted revenue.** What arrives through AI referrals?

Build the habit: prompt testing

Once a month, run the same buying questions through ChatGPT, Gemini, Claude and Perplexity. Record who gets cited, which pages get quoted, and which follow-ups the tools suggest. Those follow-ups are a free content roadmap. In our audit of 712 B2B companies, only 9% could segment AI-driven traffic at all, and 78% either never looked or saw referrals they could not attribute.¹⁴

Four myths worth retiring



“If we rank on Google, we’ll show up in ChatGPT too.”

NOT ALWAYS

ChatGPT uses Bing, not Google. Only about 11% of the most-cited domains are shared between ChatGPT and Perplexity. The engines draw on different sources.¹¹

“Schema is a nice-to-have.”

IT IS NOW ESSENTIAL

Across 24,800 analysed answers, FAQ pages with schema were cited 1.9x as often as the same content without it, and comparison pages 2.8x as often as feature pages.¹²

“AI platforms send no traffic, so ignore them.”

SOME DO, AND IT CONVERTS

Perplexity links out by design, AI-referred visitors convert at 4.4x organic⁶, and AI-referred retail traffic bounces 27% less.¹³

“AI search kills SEO.”

IT CHANGES IT

Clicks on the first result have halved in two years, and halved again where AI Overviews appear.⁴ SEO still matters. It just has a new job: getting you quoted.

ALSO TRUE


Using AI to draft is fine. Publishing it unedited is not. Sourceless claims, generic phrasing and no point of view read as low trust to readers and models alike.

We audited 712 B2B companies. 89% scored below 70.




Across SaaS, ecommerce, fintech, healthcare and services, the average AI Visibility Score is 50 out of 100. Scores come from five dimensions: platform breadth, answer accuracy, content readiness, measurement, and whether AI mentions you at all when asked about your category.¹⁴

Score distribution



Most companies are visible enough that AI knows they exist, but are not consistently cited or recommended.

Average score by industry



Ecommerce leads because product pages already use structured data and answer direct questions. Fintech and healthcare trail because compliance-heavy copy avoids plain statements.

FINDING 01

Nobody is ahead yet

Three points separate the best industry from the worst. The field is still open.

FINDING 02

Almost nobody is measuring

Only 9% can segment AI-driven traffic. 78% either don't look or can't attribute it.

FINDING 03

Small companies score higher

Companies of 11 to 50 people scored highest. Those over 1,000 scored below them.



Quick wins for the next 30 days

None of this needs a rebuild. Sequence by impact and effort, and put a name against every line.

This week

High impact, low effort

Allow GPTBot, BingBot and PerplexityBot

Verify and submit sitemaps to Bing

Add schema to your top 20 pages

Add author bios and dates to key content

This month

High impact, medium effort

Rewrite five money pages with the answer up top

Publish on two more surfaces

Split one broad page into focused pages

Start a monthly prompt-test log

This quarter

Long-term payoff

Build one tool AI cannot summarise

Create topic hubs around core buying questions

Publish original data or research

Audit agent task completion end to end

TACTIC

WHAT TO DO

WHY IT MATTERS

Entity-rich language

Name specifics: places, standards, models, materials.

Clear entities are what engines index and recall.

Purposeful internal links

Link related pages with anchor text that states the relationship.

Turns scattered pages into a readable topic network.

Answer-shaped writing

Q&As, lists, how-tos. Short, self-contained answers.

The easiest formats to lift and quote.

Topical depth

Cover a niche completely instead of one post per idea.

Depth is how trust gets inferred at scale.



The AI-readiness audit

Ten questions. Tick what applies, then count. Two or more blanks puts you with the 89% of companies that scored below 70 on page 10.¹⁴

-
- ☐ Our sitemap is submitted and verified in Bing Webmaster Tools.
 - ☐ Our key pages carry valid, complete schema markup.
 - ☐ GPTBot and PerplexityBot are allowed to crawl us.
 - ☐ Our content reads without JavaScript and loads fast on mobile.
 - ☐ Every important page opens with a direct 40 to 60 word answer.
 - ☐ Authors, dates and credentials are visible on our expert content.
 - ☐ We appear on at least two platforms beyond our own website.
 - ☐ We own an asset an AI answer cannot replace: data, tool or research.
 - ☐ An AI agent can complete a real task on our site end to end.
 - ☐ Someone owns a monthly prompt test across the three main engines.
-

SCORE

_____ of 10

FIRST THREE FIXES, AND WHO OWNS THEM

Where the numbers come from

These studies measure different things: browsing behaviour, keyword click-through, panel data. Their figures are not directly comparable. Where they disagree, we note the method. Platform share and user counts change quarterly; treat them as of mid-2026.

1	SparkToro / Datos. Zero-click study, June 2026. 68% of US Google searches ended without a click; about 45% a decade earlier.	9	OpenAI and Google. OpenAI confirmed 1bn+ weekly active ChatGPT users, July 2026. Gemini passed 1bn monthly active users, August 2026. Google reports AI Overviews at 2.5bn monthly users.
2	Similarweb. Clickstream panel, 2026. Google AI Mode refers on 1.6 to 2.5% of queries against 17 to 19% for classic results.	10	Salt Agency. September 2025. 2,138 sites: Gemini renders JavaScript reliably; Claude and Perplexity read static HTML only; ChatGPT is inconsistent.
3	Pew Research Center. July 2025. 900 US adults, 68,879 real searches: 8% click rate with an AI summary against 15% without; 1% clicked links inside the summary; 26% of sessions ended there against 16%.	11	5W AI Platform Citation Source Index. 2026. 680M citations: about 11% of top-cited domains shared between ChatGPT and Perplexity; the top 15 domains hold 68% of citation share.
4	Ahrefs. February 2026. 300,000 keywords, Search Console data: position-one click-through 58% lower where an AI Overview appears.	12	CapstonAI benchmark. Q1 2026. 24,800 responses across four engines: comparison pages cited 2.8× more than feature pages, FAQ pages with schema 1.9× more, listicles 1.7× more.
5	Seer Interactive. November 2025. 25.1M impressions across 42 organisations: organic CTR on AI Overview queries fell from 1.76% to 0.61%; cited brands saw +35% organic and +91% paid click-through against uncited brands on the same page.	13	Adobe. Retail traffic referred by AI platforms shows 27% lower bounce rates and higher conversion than traditional organic.
6	Semrush. June 2025. 500+ high-value topics: AI-referred visitors convert at 4.4× traditional organic visitors.	14	Mojo Dojo. “State of B2B AI Visibility 2026,” June 2026. 712 B2B companies across five industries, scored 0 to 100 on platform breadth, answer accuracy, content readiness, measurement and competitive coverage. Refreshed quarterly.
7	Chartbeat. 2,500+ publisher sites, December 2024 to December 2025: Google pageviews down 34% globally and 38% in the US.	15	Google Search blog. Google describes “query fan-out” as its own mechanism for dividing a search into parallel sub-queries.
8	Similarweb, via Momentic. May 2026. Worldwide assistant web-visit share: ChatGPT 53.9%, Gemini 27.9%, Claude 9.2%, Perplexity 1.3%, Copilot 1.3%. Apps add a further 15 to 60% depending on the assistant.		

Figure 1 uses illustrative similarity values to explain embeddings; it is not measured data. Figure 2 is a schematic.

Get into the answer.

AI search has changed how your market forms an opinion. Brands that get quoted early tend to keep getting quoted. Most of the 712 companies we audited could be overtaken in a single quarter of focused work.¹⁴

START HERE THIS WEEK

01

Run the worksheet on page 12 with your marketing and web owners.

02

Score yourself against the 712-company benchmark. Four minutes, ten questions.

03

Fix the two cheapest gaps before the next board meeting.

Ready to see where you stand?

Get your free AI visibility audit →

mojodojo.io/au/strategy-form

AUSTRALIA

1800 959 990

Suite 20 / Level 5
150 Albert Rd, South Melbourne VIC
3205

UNITED STATES

888 234 2349

730 Ripple Brook Ln
Elgin, IL 60120

UNITED KINGDOM

020 4634 5777

Unit 61, 38-40 The Oval
London E2 9DT

