PRACTICAL AI VISIBILITY RESOURCE

How does AI change the way customers *find and choose brands?*

Clear answers to the questions brand teams ask about AI-led discovery, research and comparison, including what shapes an answer and what it could mean for your brand.

> **Evidence note.** The question themes below are drawn from public practitioner questions about being mentioned, cited and measured in AI answers, including this SEO discussion and this OpenAI community question. These are examples of expressed demand, not search-volume rankings. These questions are not ranked by Pulse GA4 or Search Console demand data. Product-specific claims are linked to primary documentation below.

1. What counts as AI visibility?

Start with the customer decision. A brand can be **mentioned** in an answer, named first, described accurately, recommended for a specific need, or supported by a visible citation. Those are different observations. Record each separately. A citation to your site does not prove the answer recommends you; a brand mention without a citation does not tell you which source influenced it.

Example: a buyer asks for an agency that can explain a launch across Indonesia and Singapore. Note whether your brand appears, where it appears, what capability is attributed to it, which alternatives appear, and whether linked evidence actually backs the claim. Then ask whether the answer would help a buyer make a sound choice.

2. How can our brand appear in ChatGPT, Google AI answers or Copilot?

Think across the customer journey: discovering options, researching claims and comparing providers before a decision. Make the information a person needs easy to find, understand and verify. Publish clear pages about who you serve, what you do, where you operate, and the evidence for important claims. Keep facts consistent across your own site and credible independent sources. For Google AI features, Google says established Search practices remain relevant; it does not require special AI markup. For ChatGPT search, OpenAI documents OAI-SearchBot and separate crawler controls. None of these steps guarantees inclusion or a particular answer.

3. Why does a competitor appear when we do not?

Diagnose the specific question before changing content. Your page may not answer that customer need, the engine may not retrieve it, the answer may rely on other sources, or your public proof may be thinner for that market. A technical block is

another possibility. Compare the answer, cited URLs and page content with your competitor's. Do not infer a universal ranking from one prompt.

Make a gap sheet with columns for customer question, engine, market, language, answer date, brands named, first mention, cited sources, and the claim to investigate. Check a small set of real questions repeatedly. If the answer changes, preserve both observations.

4. How do we measure AI visibility without pretending to know every prompt?

Use two layers. First, take an **observed answer sample**: a documented set of customer questions, engines, markets and languages, with dates and the exact answer and sources. Second, use available platform data for discovery and referrals. Bing Webmaster Tools AI Performance reports citations across supported Microsoft AI experiences, but its aggregates do not show the role or placement of each cited page. Search Console queries and page metrics can add context, while some low-volume queries are anonymized. Label what is measured, sampled or inferred.

For any rate, show the count and denominator in a defined sample of customer questions, engines, markets and dates. A real example comes from the PRecious Banking AI Authority Index: in 706 of 1,628 eligible matched comparisons, two tested AI service routes answering the same banking question named a different bank first (43%). The 1,628 is the number of route pairs where both answers named a recognised bank; 706 of those pairs disagreed on the first bank. Pairs can share an underlying answer, so these are not 1,628 independent tests. First mention describes answer order, not endorsement, trust or customer preference. This August 2026 banking sample does not estimate a rate for other brands or industries. For your own brand, record whether it appears and whether its description is accurate; count visible citations separately. Keep platform impressions, site visits from AI referrals and qualified leads separate when available. A citation, visit or lead cannot be inferred from a mention alone.

5. Why is one AI visibility score misleading?

Answers vary with the question, engine, source conditions, location, language and time. A broad average can conceal a critical buyer question where your brand is absent or misdescribed. The PRecious Banking AI Authority Index found different first-mentioned banks in 43% of eligible matched route comparisons for the same question. That finding is about its tested banking sample, not a rate for every brand or industry. For your brand, segment the sample by the decisions that matter and report coverage and uncertainty beside any summary.

6. What sources influence an AI answer?

Inspect the sources the answer actually presents. They may include your pages, independent journalism, customer reviews, official registers, research or comparison pages. A link in an answer shows a source was presented, not that every statement came from it. Open the source and test whether it supports the specific claim. Ask what a customer would trust if they clicked through.

For a disputed claim, make a source map: statement, visible source, supporting passage, owner, correction needed and date checked. The communications response might be a clearer product page, an expert explanation, an updated public

fact, or credible outside validation. Avoid manufacturing reviews or citations.

7. What content should we make first?

Start with unanswered customer questions, not a list of fashionable AI keywords. A useful page names the audience and situation, gives a direct answer, explains limits, provides current proof, and offers a next step. Comparison pages can be valuable when they are fair and specific. For a B2B service, explain the job to be done, where it fits, what it does not cover, what evidence supports outcomes, and how someone can evaluate alternatives.

Choose one page to improve using a real sales or support question. Have an expert check every claim, add the supporting source, update the date, then retest the same questions. Google's guidance emphasizes helpful content and warns against overfocusing on special markup.

8. What should the web team check?

Confirm that important pages are reachable, indexable where intended, linked from the site, and readable without a login or script-only rendering. Inspect robots.txt, noindex directives, canonical URLs, HTTP status and CDN or firewall rules. For ChatGPT search specifically, distinguish OAI-SearchBot from GPTBot; their purposes and controls differ. A crawler allowance does not guarantee a citation. Ask the web team for a short evidence record: URL, status, indexing/crawl result, last check and any blocked route.

9. Will FAQ schema or an llms.txt file make AI recommend us?

No published platform rule promises that. Structured data can help describe content when it matches the visible page, but Google says no special schema is required for its AI features. Google also restricted FAQ rich results primarily to well-known government and health sites. Use schema for a supported, truthful purpose and test it. Treat llms.txt as an optional experiment, not a substitute for crawlable, useful pages or public proof.

10. Why test language and market separately?

Translation changes the wording, while market context changes the useful examples, regulations, competitors and sources. Test natural local questions, not only word-for-word translations. Record the engine settings and query date. Have a local reviewer judge whether a description would be helpful and accurate for that audience. If the answer differs, inspect the sources and page coverage before deciding what to change.

11. What can communications teams actually influence?

Communications can improve the quality and consistency of public evidence: clear owned pages, accurate spokesperson explanations, independent expert commentary, customer proof where rights allow, and corrections to outdated facts. Build one claim map across web, PR, sales and social content. Give each claim an owner, source, market and review date. Independent coverage should earn its place through relevance and accuracy, not a promise of AI citation. If a placement is

published, test whether it is actually presented as a source in repeated answers before attributing any change to it.

Separate leading indicators (pages improved, errors corrected, sources validated) from observed answer changes and business outcomes. Retesting may show no change, a better description, or a new gap. All are useful evidence.

12. What should we do in the next 30 days?

1. **Week 1:** collect 10 real customer questions from sales, search and support; select engines, markets and languages.
2. **Week 2:** record answers, first mentions, descriptions and cited URLs. Check source-to-claim fit.
3. **Week 3:** fix one technical issue and improve one high-value answer page with verified facts and local context.
4. **Week 4:** rerun the same questions, note changes and uncertainty, then decide the next content or communications action.

For a quick team starting point, use the free AI Visibility Readiness Check. It assesses how you check AI answers and act on gaps in website content, public proof and communications. It does not measure live visibility. For a measured review of your brand, ask PRecious.

Sources and further reading

- Google Search Central: AI features and your website
- Google Search Central: optimizing for generative AI features
- OpenAI: crawler controls
- OpenAI: publishers and developers FAQ
- Bing Webmaster Tools: AI Performance
- Google Search Console: data limits

Next: read the plain-language glossary.