RESEARCH BRIEF 03 / OCTOBER 2026

Measuring AI search visibility

A 60-day observation protocol

Define the questions, conditions and counting rules before changing a page. Track mentions, citations and commercial outcomes in separate records.

IN THIS BRIEF

01 Define a repeatable question panel
02 Count and interpret observations
03 Keep a reusable measurement log

A brand appears in one AI-generated answer. After a page update, it appears again. That is an observation, but it does not yet show whether visibility improved or whether the update caused anything.

A useful measurement plan defines the questions, products, conditions and counting rules before the first page edit. This article proposes a small operational panel a business can maintain over sixty days. It is not an official platform metric or a representative survey of all customer searches.

1. Define the market and search task

Choose one market, language and offer. For example, an English-language service aimed at UAE business owners is one panel. Arabic questions require their own wording and separate reporting; translating the English set does not establish Arabic customer demand.

Select ten non-branded questions relevant to discovery and evaluation. Include questions about fit, scope and buying risk. Keep brand-navigation questions in a separate panel so asking for your own company does not inflate discovery results.

Record why each question was chosen. Search Console queries, documented sales questions and customer research can support the choice. A team brainstorming session is a hypothesis, not demand evidence.

2. Separate technical eligibility from observed inclusion

For Google AI Overviews and AI Mode, Google describes ordinary search eligibility: a supporting page must be indexed and eligible for a search snippet. Its guidance does not require special AI files or additional schema for those features. These are Google-specific statements, not universal rules for every assistant.

Source: [1] Google Search: AI features and measurement

Check each platform's access policy separately. For example, OpenAI documents OAI-SearchBot for search and GPTBot for potential model-training use, with independent controls. An access change is not itself evidence of a mention or citation.

Source: [2] OpenAI: search and training crawlers

3. Freeze the observation method

At day 0, 14, 30 and 60, run each of the ten questions three times per selected product in fresh conversations or equivalent independent sessions where available. Record any inability to isolate sessions. Repetition describes variability within this panel; it does not make the observations statistically independent.

Keep exact wording, language, account state, search setting, observed location and collection window consistent where possible. Record the product and visible model/version, or unknown when it is not exposed. An API run belongs in its own dataset unless its equivalence to the consumer interface is established.

The proposed design creates thirty scheduled observations per product per round. Save errors and missing observations separately. A timeout is not a valid answer that omitted your brand. In Google Search, also distinguish an AI feature that did not appear from an AI answer that appeared without mentioning you.

4. Count mentions and citations separately

Measure	Counting rule
Brand mention rate	Valid generated answers naming the intended business, divided by valid generated answers inspected
Owned-domain citation rate	Valid generated answers with at least one verified citation to your domain, divided by valid generated answers inspected
Citation concentration	Citation-bearing answers grouped by the actual destination page
Accuracy notes	Manually recorded factual errors, with the answer and verification source

Count an answer once for each rate, even if it mentions the brand repeatedly. Check ambiguous company names. Record the cited destination rather than assuming a visible label or a third-party mention links to your site.

Keep raw counts beside percentages. If the valid-answer count falls, show that change rather than silently comparing percentages with different coverage.

5. Use the results to choose a page-level action

Suppose, in an illustrative dataset, owned-domain citations increase from six of thirty valid answers to nine of thirty. The panel citation rate moves from 20% to 30%, an increase of ten percentage points. This does not establish a market-wide share or a causal SEO improvement. The questions were selected rather than randomly sampled, and repeated outputs can be related.

Inspect which questions changed and which pages were cited. Three extra citations may come from one question while the other nine remain unchanged. That pattern suggests a narrower next investigation than a broad claim that visibility increased everywhere.

Keep a change log: URL, edit, rationale, release date and next observation window. Avoid changing the prompt set halfway through the comparison. If the product changes its model or interface, mark a break in the series. Unchanged comparison pages can add context, but they do not create a controlled experiment by themselves.

6. Keep search and commercial outcomes in separate records

Google says its AI search features are included in Search Console's Web performance data. A change in that overall report should not automatically be labelled AI-only growth.

Source: [1] Google Search: AI features and measurement

Review referral traffic where available, enquiries with evidence and qualified opportunities separately from prompt observations. Missing referral data does not prove that no AI-assisted visit occurred. A citation does not prove a sale.

Keep this observation log

Field	What to record
Date	Observation date and collection window.
product/version	Product and visible version; unknown if unavailable.
market/language	Intended market, language and observed location.
question ID	Stable identifier for the comparison.
exact prompt	The question wording used in this run.
session conditions	Account state, search setting and known limitations.
feature present	Separate no generated feature from no brand mention.
valid answer	Flag errors and missing observations separately.
brand mention	Confirm the intended business identity.
cited destination	Actual linked page, not just its visible label.
accuracy note	Observed error and verification source.
evidence	Saved answer or permitted capture.

At day 60, report what changed in this defined panel, which pages and questions contributed, what else changed during the period and which commercial outcomes can actually be traced. That is enough to support a useful next decision without pretending the panel represents the whole market.

The panel size, schedule, worksheet and rates are Layerly proposals. All example counts are illustrative. No measured Layerly visibility increase or client result is asserted.

Technical sources

Primary documentation reviewed 30 October 2026. The diagnostic frameworks and illustrative examples are identified in the article.

[1] Google Search: AI features and measurement
https://developers.google.com/search/docs/appearance/ai-features

[2] OpenAI: search and training crawlers
https://developers.openai.com/api/docs/bots

layerly