

GOOGLE AIO · AI VISIBILITY · CONTENT STRATEGY · B2B STRATEGY · SEO

Getting cited in Google AI Overviews: what B2B brands need to know

Google AI Overviews now trigger on 82% of B2B Technology queries. The ranking signals that drive citation are different from standard SEO: top-10 position is no longer a gate, JSON-LD schema may reduce citations for well-performing pages, and query fan-out means topical breadth matters more than keyword depth. This article sets out what the controlled evidence shows and where the structural limits lie.

AIVIARA RESEARCH · MAY 2026

If your category queries are generating AI Overviews (AIO) and your brand is not appearing in them, you are invisible at exactly the moment many buyers are forming their evaluation shortlist. That gap is widening quickly.

82% of B2B Technology queries triggered Google AI Overviews by December 2025, up from 36% in May 2025.

SEARCH ENGINE JOURNAL ANALYSIS · DECEMBER 2025

The public evidence around AI Overviews is fragmented, contradictory in places, and changing quickly. But several patterns are now consistent enough to act on. Two findings are worth flagging before the detail, because they contradict what most guidance in this space says.

First, organic ranking still matters for AIO citation, but far less than it did a year ago. Second, adding JSON-LD schema markup may not help and, in a controlled study of 1,885 pages, produced a statistically significant decline in citations for pages already receiving them.

Your ranking position is still relevant, but it's no longer the gate

Most AEO guidance treats AIO citation like featured snippets: rank in the top five and your content gets included. That assumption has become unreliable.

Ahrefs analysed 863,000 keywords and 4 million AIO URLs in March 2026. They found 38% of cited pages ranked in Google's top 10 for the same query, down from 76% in their July 2025 study. The remaining citations split almost evenly between positions 11 to 100 (31.2%) and beyond position 100 (31.0%). Ahrefs notes their own methodology caveat: improved parsing in the 2026 study may have caught citations from outside the top 10 that earlier methods missed. The direction of travel is still clear, even accounting for that.

Only **38%** of AIO-cited pages ranked in Google's top 10 for the same query. Down from **76%** in July 2025.

AHREFS — 863,000 KEYWORDS, 4 MILLION AIO URLs · MARCH 2026

Originality.AI's November 2025 study put numbers on the other side of the picture. Of all AIO citations, 48% overlapped with the top-100 organic results. The other 52% came from sources outside the top 100 entirely. Within the 48% that did overlap with organic rankings, 52% came from top-10 pages (which works out to roughly 25% of all AIO citations). The study's headline, often misread, is that the majority of AIO citations do not come from pages ranking in the top 10.

seoClarity's October 2025 analysis of 362,000 desktop queries that triggered AIOs found stronger top-10 correlation: position 1 URLs were cited 43% of the time, position 20 URLs 7% of the time. Top-20 organic results had at least one URL in the AIO for 94% of queries.

Ranking in the top 20 still materially increases citation probability. It is just no longer a gate. And the mechanism that has weakened the correlation is important enough to understand directly.

How Google's query fan-out changes the content brief

Google has explicitly described query fan-out in AI Mode documentation: rather than running a single search and picking the top results, the system breaks the original query into multiple sub-queries, each exploring a different facet of the question, and retrieves sources for each. Aleyda Solis describes this as "decomposing the query into subtopics and issuing a multitude of queries simultaneously." Researchers believe similar retrieval behaviour operates in standard AI Overviews, though Google has not confirmed this explicitly.

If fan-out applies to standard AIO as researchers infer, the practical consequence for content strategy is significant: a page that ranks #1 for the head term but covers only one

angle may be passed over in favour of pages that answer more of those sub-queries. Covering the topic from multiple angles matters more than deep optimisation for a single keyword.

Consider a buyer searching "best workflow automation tool for operations teams." Fan-out might generate sub-queries around integration capabilities, pricing models, team adoption, technical requirements, and migration from legacy tools. A page that ranks well for "workflow automation software" but covers only product features is competing against pages that address the full evaluation picture. The evaluation query is the one with the AIO; the narrow page may rank well but get cited for none of the sub-queries.

What answer-first structure actually means

Google's AI Overview cites passages, not pages. The system identifies specific sections that directly answer sub-query questions. Content buried under background framing and scene-setting is harder for retrieval systems to extract.

Search Engine Land characterises AIO inclusion as "a retrieval problem, not a ranking issue." The page may rank well, but if the relevant passage is not structured for extraction, the page gets skipped.

Here is what this looks like in practice.

Before:

"Vendor risk management has become a strategic priority for operations leaders as supply chains grow more complex and regulatory requirements intensify. In this environment, organisations are increasingly looking for frameworks that..."

After:

"Vendor risk scoring models evaluate suppliers across four dimensions: financial stability, security posture, compliance exposure, and operational continuity. A score below 60 on any dimension typically triggers a review before contract renewal."

The second version answers a question. The first does not answer anything; it contextualises. Retrieval systems extract the version that answers; the contextualising version gets passed over even if the page around it ranks well.

The formatting rules follow from this. Lead every major section with the direct answer to the question the section addresses. Use H2 and H3 headers that reflect how a buyer phrases the query, not how a copywriter labels a section. Supporting context and caveats come after the answer, not before it.

The schema finding that contradicts most advice

Adding FAQ schema and structured data markup is standard AEO advice. An Ahrefs controlled study published in 2026 tested it directly, and the results are not what most guidance suggests.

Ahrefs tracked 1,885 pages that added JSON-LD schema markup between August 2025 and March 2026, against 4,000 matched control pages. For Google AI Overviews specifically, the treated pages showed a statistically significant decline of approximately 12 daily citations per affected page — measurable, and in the wrong direction. For AI Mode and ChatGPT, results were close to zero. Ahrefs concluded that for pages already performing well in AIOs, JSON-LD schema is not the lever that drives citation.

Pages that added JSON-LD schema saw a statistically significant **decline of ~12 daily citations** in Google AI Overviews — measurable, and in the wrong direction.

AHREFS CONTROLLED STUDY — 1,885 TREATED PAGES, 4,000 CONTROLS · 2026

There is a correlation, but it runs in the opposite direction from what the controlled evidence supports. Ahrefs found that pages appearing in AIOs are roughly three times more likely to have schema than pages that don't. The available controlled evidence does not show schema independently increasing citation likelihood for pages already performing strongly in AIOs; well-optimised pages tend to have both schema and citation, because they are better pages overall.

One important scope note on this finding: the Ahrefs study tested pages already receiving 100 or more AIO citations (pages that were already being selected regularly). Whether schema helps pages with no existing AIO presence gain their first citations is an open question the study cannot answer. If your content is not currently being cited at all, the schema finding may not apply in the same direction. What it does tell you is that schema is not the primary intervention to prioritise if you are already appearing.

The query types that matter for B2B

Not all query types trigger AI Overviews equally, and the pattern is commercially significant for B2B content strategy.

Question-format queries trigger AIOs 57.9% of the time, versus 15.5% for non-question formats.

ZIPTIE ANALYSIS OF SEARCH ENGINE JOURNAL DATA

ZipTie's analysis found question-format queries trigger AIOs 57.9% of the time, versus 15.5% for non-question formats. Discovered Labs' analysis found seven-word-plus queries trigger AIOs 46.4% of the time, versus 9.5% for single-word searches. B2B evaluation queries are naturally both question-shaped and specific ("how does [product] integrate with [system]," "what's the best [category] for [use case] with [requirement]"), meaning AIO triggers far more often for B2B research behaviour than for broad consumer searches.

Comparison queries have grown dramatically. According to Discovered Labs, "best [product]" queries grew from 5% AIO presence around Black Friday 2024 to 83% by Black Friday 2025. Vendor comparison is where B2B buyers spend time, and that query type now almost always generates an AIO.

Two patterns are notably absent from AIO:

Branded queries. Only 4.79% of branded keywords triggered AIOs, according to analysis cited by Discovered Labs. Navigation-intent queries (a buyer going directly to a vendor's website) are systematically excluded. You do not need to optimise your brand homepage for AIO inclusion.

High-CPC queries. Discovered Labs' analysis found keywords above \$10 CPC triggering AIOs roughly 17% of the time, against 60-plus percent for keywords below \$0.50 CPC. The mechanism is not confirmed, but the correlation is strong enough to matter: the bottom-funnel purchase-intent queries where B2B brands most want to appear may be systematically less likely to trigger an AIO in the first place. AIO is more active at the evaluation and consideration stage than at the final-decision stage.

Two failure modes worth knowing before you start

The wrong question version. Search Engine Land identifies this as the primary reason content that ranks well still fails to appear in AIOs. With query fan-out generating multiple sub-queries, your content may be optimised for the head term but not for the specific

sub-question AIO is actually trying to answer. The fix is not keyword optimisation in the traditional sense — it is auditing your content against the full set of questions a buyer might have around the topic, not just the query you are targeting.

A practical example: a software vendor's page on "enterprise data governance" might rank for that term and be well-written. But if AIO is answering sub-queries like "how long does data governance implementation take" and "what compliance frameworks does [category] support," and those specific questions are not addressed on the page, the page will not be selected for those sub-queries regardless of its ranking position.

The promotional content filter. Inspired Marketing B2B describes AI systems as "conservative by design." Content that reads as promotional rather than informational is filtered out of AIO selection. This is a selection mechanism based on content purpose, not a content quality judgment. Pages written to convert (feature lists with benefit statements, product-first framing, free-trial CTAs mid-page) are oriented toward a different job than answer provision. Practitioners consistently report that heavily promotional content is less likely to be selected for AIO inclusion, regardless of where it ranks.

The test is whether a section can stand alone as an answer to a question. A product page section that reads "Our advanced workflow engine handles 10,000 operations per second" answers a question. A section that reads "Experience the power of our industry-leading workflow automation platform" answers nothing and is not extractable.

Start with technical access

Before content structure has any effect, Google needs to be able to reach and index your pages. The same crawl and indexation requirements that apply to standard organic search apply here, so it is worth checking explicitly before investing in AIO content work.

Google's documentation states that AIO eligibility follows standard Search eligibility: being indexed and eligible for snippets. If your pages are blocked from crawling or excluded from the index, they cannot appear in AIOs regardless of content quality.

Some enterprise CMS and hosting configurations include broad crawler disallow rules that may inadvertently block Googlebot, often added during security hardening without specific consideration for search eligibility. Check your robots.txt against Google's current crawler user agents before assuming access is clean.

Google's official guidance states directly that "There are no additional requirements to appear in AI Overviews or AI Mode, nor other special optimizations necessary." AIO

eligibility is not behind a separate application or approval; it follows from standard indexation.

Where to start, and in what order

These actions are sequenced by expected impact relative to effort.

Map your evaluation-stage query set first. Identify the queries your buyers use at the vendor comparison stage, not the branded or bottom-funnel queries. Question-format queries ("how to," "what's the best," "[product A] vs [product B]"), category definition queries, and use-case-specific queries are the AIO-active population for B2B. Use these as your target set for citation monitoring and content auditing.

The highest-leverage structural change: answer-first rewrites. For each page in the target set, check whether each major section opens with a direct answer to the question the heading implies. If a section leads with background context, restructure it. Lead with the answer; support with evidence.

Before adding AEO changes to existing pages, check whether those pages currently rank in Google's top 5 for their primary keywords. For pages with strong organic performance, add answer-first intros and FAQ blocks as additive elements rather than rewriting existing sections. Rewriting a page that generates consistent organic traffic to optimise for AIO citation is a trade that needs to be made with data, not as a default move.

Topical depth across the content library. With query fan-out active, individual highly-optimised pages are less important than a content library that covers the topic from multiple angles. Audit whether your content addresses the full range of questions a buyer has at the evaluation stage: integration, pricing models, implementation, compliance, migration from alternatives, category comparisons. Each angle that is missing is a sub-query where you cannot be selected.

What usually gets missed: standalone FAQ and definition pages. Q&A; formats and direct-answer content are the highest-performing citation formats across multiple analyses. Dedicated pages that answer specific category questions (not product questions) are underrepresented in most B2B content libraries. A page that directly answers "how does [category] handle [specific requirement]" can appear in AIOs for that sub-query even if it does not have strong overall domain authority.

Entity consistency across the web. Inconsistent brand descriptions, category labels, and product terminology across different channels create ambiguity in Google's entity graph. If

Google cannot confidently resolve what your brand is and what it does, it tends to skip it in favour of sources it can. Ensure your brand's description, category, and key capabilities are consistently stated across your website, press coverage, and third-party mentions. This is practitioner inference rather than a controlled finding, but the mechanism is plausible and the work is low-cost.

How to know if it's working

Standard analytics will not give you a reliable read on AIO citation presence. Google does not pass citation attribution in a way that surfaces cleanly in web analytics. You may see some referral traffic from AI Overviews but it significantly undercounts.

The more reliable method is direct query monitoring. Run the queries in your evaluation-stage set in Google search, record whether an AIO appears, and record whether your content is cited. Do this weekly for a set of 20 to 30 tracked queries. It is manual and impractical at scale, but it establishes a baseline that analytics cannot.

If you are tracking organic traffic to pages where AIO is active for competitive queries, you may see traffic declines that reflect AIO's presence rather than a content problem. Understanding the two effects separately requires query-level monitoring, not just traffic reporting.

Set a 60 to 90 day window before expecting measurable movement from content changes. Pages need to be re-crawled and re-indexed before AIO selection changes. Schema changes may be picked up faster but, as the Ahrefs study shows, schema is not the primary lever here.

What this playbook won't achieve

Domain concentration is a structural ceiling. According to The Digital Bloom's 2025 practitioner analysis of citation patterns (read as industry observation rather than independent research), the top 10 domains account for 53.87% of all AIO citations, and the top 20 take 66.18%. Wikipedia takes 11.22% of all citations; YouTube takes 9.51%. Google's own properties appear in 43% of all AI Overviews and account for 22.81% of citations collectively. There is no equivalent of a neutral authority aggregator for B2B tech in that concentration tier. B2B brands can and do appear in AIOs, but citation share at the category level is constrained by sources no amount of content optimisation will displace.

High-CPC queries are systematically less active for AIO. As noted in Discovered Labs' analysis, bottom-funnel purchase-intent keywords (typically high-CPC) trigger AIOs at substantially lower rates than the evaluation and educational queries where AIO is most active. The queries where AIO is most active for B2B (evaluation, comparison, educational) are earlier in the buying process. AIO citation builds brand presence and shortlist consideration. It is not reliably a bottom-funnel conversion mechanism.

Branded queries are excluded by design. 4.79% of branded keyword queries trigger AIOs. Optimising for AIO citation on queries where buyers are already searching for you by name is not a productive use of effort. The citation opportunity is in category and evaluation queries, where you are competing for attention before the buyer has formed a vendor preference.

AI Mode is a different surface with different behaviour. AI Mode, which became a standard mobile feature in December 2025, generates citations independently from standard AIO results. Google has confirmed that AI Mode uses query fan-out extensively, but how its citation logic differs from standard AIO is not publicly documented. Brands appearing consistently in standard AIOs may see different rates in AI Mode. Monitoring both surfaces separately is the right approach, and the expectation should be that the optimisation rules may diverge as the surface matures.

The measurement gap

Across all the data in this article, one question remains genuinely open: whether the weakening correlation between top-10 ranking and AIO citation reflects a change in Google's weighting of organic rank as a signal, or simply an expansion in available citation slots (average citation count grew from 6.82 per AIO in November 2024 to 15.22 by February 2026, according to SE Ranking's analysis). More slots means content from outside the top 10 gets included simply because the top 10 cannot fill them all. If that is the mechanism, the prescription is different from what it would be if Google is actively deprioritising ranking as a selection signal.

Aiviara is building infrastructure for monitoring AI brand citations and factual accuracy across LLM platforms. Early access information is available at aiviara.com.