

AI Search Visibility Guide for B2B Marketers

WHAT YOU NEED TO KNOW
TO INFLUENCE BRAND VISIBILITY

01

Actual AI Usage Data (Finally!)
to Compare Against Google

4

02

A Peek at How ChatGPT Pulls
and Prioritizes Information

7

03

A Detailed Look at the Hidden
Fan-Out Queries Behind Most
AI Answers

9

04

A Deep Dive into How ChatGPT
Scores Content After Retrieval

14

05

The Rules of AI-Readable
Content: Structure, Clarity
& Segmentability

21

06

How PR & Off-Site Signals
Quietly Shape AI Visibility

25

07

How Google’s AI Overviews
Work & What We Can See
Behind the Scenes

29

08

Why “Best Of” Lists Dominate
AI Answers & What To Do About It

35

09

Message Drift: How AI
Platforms Blend Your Brand
Story Across the Web

38

10

How We Can Start Measuring
& Optimizing AI Visibility Today

41

11

What All of This Means
Moving Forward

47

For years, B2B marketers understood how to play the search visibility game, from optimizing for keywords to building backlinks to tracking Google rankings. But as we look ahead, a new force is rewriting the playbook. AI search platforms like ChatGPT and Google’s AI Overviews have the potential to fundamentally change how B2B buyers find information and make decisions.

Unlike traditional search engines, these AI ecosystems can feel like a black box right now. They lack the transparency we’re used to as marketers, leaving many wondering:

- ▶ How does ChatGPT actually decide which brands to surface in its answers?
- ▶ How does it score my content?
- ▶ And perhaps most importantly, how do I ensure our brand story remains accurate when AI platforms blend information from across the entire web?

We’re here to help peel back the layers.

In this guide, we’ll bring clarity to the AI search process so you can better understand how these platforms pull, prioritize, and score information. We’ll share a behind-the-scenes look at the processes these platforms use to make decisions, along with insights to help you adapt your digital strategy and position your brand for increased visibility.



Actual AI Usage Data *(Finally!)* to Compare Against Google



When it comes to understanding what people search for, Google has always been an open book.

Tools like the Google Keyword Planner, Semrush, and industry-standard clickstream datasets give us reliable directional insight into keywords and monthly search volumes:

- ▶ What people search for
- ▶ How often they search for it
- ▶ How those trends shift over time

This transparency is what makes SEO measurable and predictable.


But when it comes to ChatGPT and other AI platforms, we’ve never had anything close to that transparency — no keyword volumes, no prompt trends, no search categories, and no data showing whether people were actually using it to find information versus using it to create content, write code, analyze data, brainstorm ideas or complete other tasks.

For the last couple of years, the constant hype around “AI replacing search” has been largely speculation, because we lacked any true activity data from AI platforms themselves or credible third-party studies. That’s finally changed. We now have a few sources of real usage data that let us see how people are behaving across search engines and AI tools.

The [State of Search](#) report from Datos and SparkToro details clickstream data from a panel of tens of millions of anonymized desktop users and the websites that they visit.

As of June 2026, traditional search still dominates total activity in the U.S., with Google holding 94% of desktop search.

AI tools, while growing steadily, made up roughly 1.83% of total desktop visits in Q2 2026, up from 1.37% a year earlier.



As of June 2026, ChatGPT is the clear leader among AI platforms, used by 36% of desktop users. Gemini is the fastest riser at 19% and Claude third at 11%, while Perplexity, DeepSeek, and Copilot cluster in the low single digits.

These figures count unique visitors per platform independently, so they can overlap — a single person using multiple tools is counted in each.

The clickstream data tells us how many people are using these tools, but to understand what they’re doing inside the leading platform, we can turn to ChatGPT directly. In September 2025, OpenAI + Harvard teamed up and released a technical whitepaper called [How People Use ChatGPT](#). The 63-page report details data through July 2025. On average:

- ▶ 2.5 billion daily prompts
- ▶ Of these, only ~21% are “seeking information”

That gives us our first real metric: ~66 million search intent prompts per day.

Now compare that to Google with ~14 billion searches per day. As a result:


ChatGPT search-like behavior is about 210 times smaller than Google — smaller than search activity on DuckDuckGo.

While far from being a “Google replacement” right now, ChatGPT is becoming an increasingly meaningful channel where audiences look for information. In July 2025, ChatGPT had 700 million weekly active users. By early 2026 that figure had [passed 900 million](#), so today’s search intent volume is almost certainly meaningfully higher.

While this guide focuses on ChatGPT as the dominant platform, the other major models like Claude, Gemini and Perplexity operate in fundamentally similar ways: retrieving, scoring and synthesizing information rather than returning a list of links.

So, the real question for marketers becomes: how does ChatGPT decide which brands and pages to surface in its answers?

Search, AI, and E-Commerce platform visits in the US




Monthly Events Share (%) = Monthly visits to domain's webpages / Total monthly webpage visits across entire panel * 100%



Source: Datos



Share of desktop users using AI in the US



Monthly % of userbase = Number of unique users visiting the domain's webpages (monthly) at least once / Total number of unique panel users (monthly) * 100%



Source: Datos



A Peek at How ChatGPT Pulls and Prioritizes Information

02

Now that we know how large (or small) ChatGPT’s search-like behavior really is, the next factor is the one that matters for all of us: *How ChatGPT decides which brands and pages to surface in its answers.*

Unlike Google, there is no ranking page, no keyword position and no click-through-rate (CTR) data to review. ChatGPT works differently, but understanding its process helps explain why some brands show up and others don’t.

Here’s the high-level overview of how ChatGPT “finds” information:

It starts with what it already knows

Although OpenAI hasn’t been fully transparent at a granular level, ChatGPT is [trained on a massive snapshot of the public web](#) that includes articles, documentation, reference material, forums, books, online conversations, and more.

When your question doesn’t require fresh details, niche specifics, breaking news or emerging information, the model answers directly from this training data. This covers a portion of everyday informational queries that it can answer fast, like: “What’s the best recipe for a chocolate birthday cake?”

When it needs newer or more specific information, it calls in Google + Bing

If ChatGPT isn’t confident that it has enough information internally on its model, it performs a search.

Contrary to popular belief, ChatGPT makes a handful of hidden searches via an API to Google and Bing, grabbing the top search results and bringing the URLs and metadata back for further analysis.

This is why SEO visibility becomes AI visibility. If you’re not showing up in search engines, ChatGPT simply can’t retrieve you.

From the URLs it retrieves, ChatGPT evaluates several broad signals

Unlike a search engine, ChatGPT isn’t ranking pages — it’s interpreting them. Once the model pulls in results, it decides which information to trust, incorporate and cite by looking for cues:

- ▶ **Clarity & structure**
Clear titles, headers and metadata help AI understand the topic instantly.
- ▶ **Relevance to the question**
Does the page directly address what the user is asking?
- ▶ **Authority & trustworthiness**
Pages from reputable sources or information reinforced across multiple sites tend to carry more weight.
- ▶ **Consistency across the web**
AI favors information that multiple independent sources agree on.
- ▶ **Freshness markers**
AI infers how current a page might be based on timestamps, language patterns and contextual clues from search engines.

These signals don’t produce a “rank,” but they strongly influence what gets blended into the final response. (More details on influencing factors later.)

The end result is a synthesized answer, not a list of links

Unlike Google, ChatGPT isn’t trying to send traffic anywhere. Its goal is to distill the most reliable information it can find — across both its training data and real-time sources — into a single coherent answer.

As a result, even highly ranked content sometimes doesn’t appear in AI outputs: ChatGPT pulls what it understands, not what simply ranks.

Now that we know what ChatGPT looks for, the next question becomes:

How does ChatGPT actually run these searches — and how does it choose which URLs to analyze or cite?

A Detailed Look at the Hidden Fan-Out Queries Behind Most AI Answers

03

Let’s look at how ChatGPT runs those searches — known as fan-out queries — and how it chooses which URLs to analyze or cite.

The platform doesn’t perform a single search the way a human does in Google or Bing.

ChatGPT runs something called a fan-out process, a behind-the-scenes method that generates multiple search queries at once. It gathers a broad pool of URLs from the answers and then narrows them down before synthesizing an answer.

Here’s how that works:

ChatGPT doesn’t run ONE search — it runs SEVERAL at once

If you ask: “Who are the leading suppliers of industrial HVAC systems?” ChatGPT won’t run that one query. It will generate a cluster of related searches, such as:

- ▶ “top industrial HVAC suppliers”
- ▶ “HVAC systems for manufacturing”
- ▶ “industrial ventilation manufacturers”

These are all fired in parallel through its search API connections.


Why multiple? Because each phrasing returns a slightly different set of URLs, giving ChatGPT a broader, more complete picture of the topic. This is the first gate where brand visibility is either won or lost. If your content doesn’t appear in these fan-out queries, your brand never enters the AI answer pool.

It pulls in the top URLs from those different queries

Across these variations, ChatGPT retrieves:

- ▶ Title tags
- ▶ Meta descriptions
- ▶ URLs
- ▶ Headings
- ▶ Small text snippets

This becomes its candidate set, or the list of pages ChatGPT will consider during synthesis. Sometimes this pool is small (10–20 URLs). Sometimes it’s large (hundreds).



ChatGPT then clusters, compares and filters the results

From that pool, ChatGPT begins analyzing:

- ▶ Which pages reinforce one another
- ▶ Which details remain consistent across sources
- ▶ Which information appears trustworthy
- ▶ Which outliers look unreliable
- ▶ Which patterns emerge across multiple URLs

ChatGPT is looking for consensus rather than just returning a list of URLs. The more your brand appears across these retrieved URLs (your own site + third-party mentions + media coverage), the more likely you are to be included in the answer.

It synthesizes the highest-confidence information into a single response

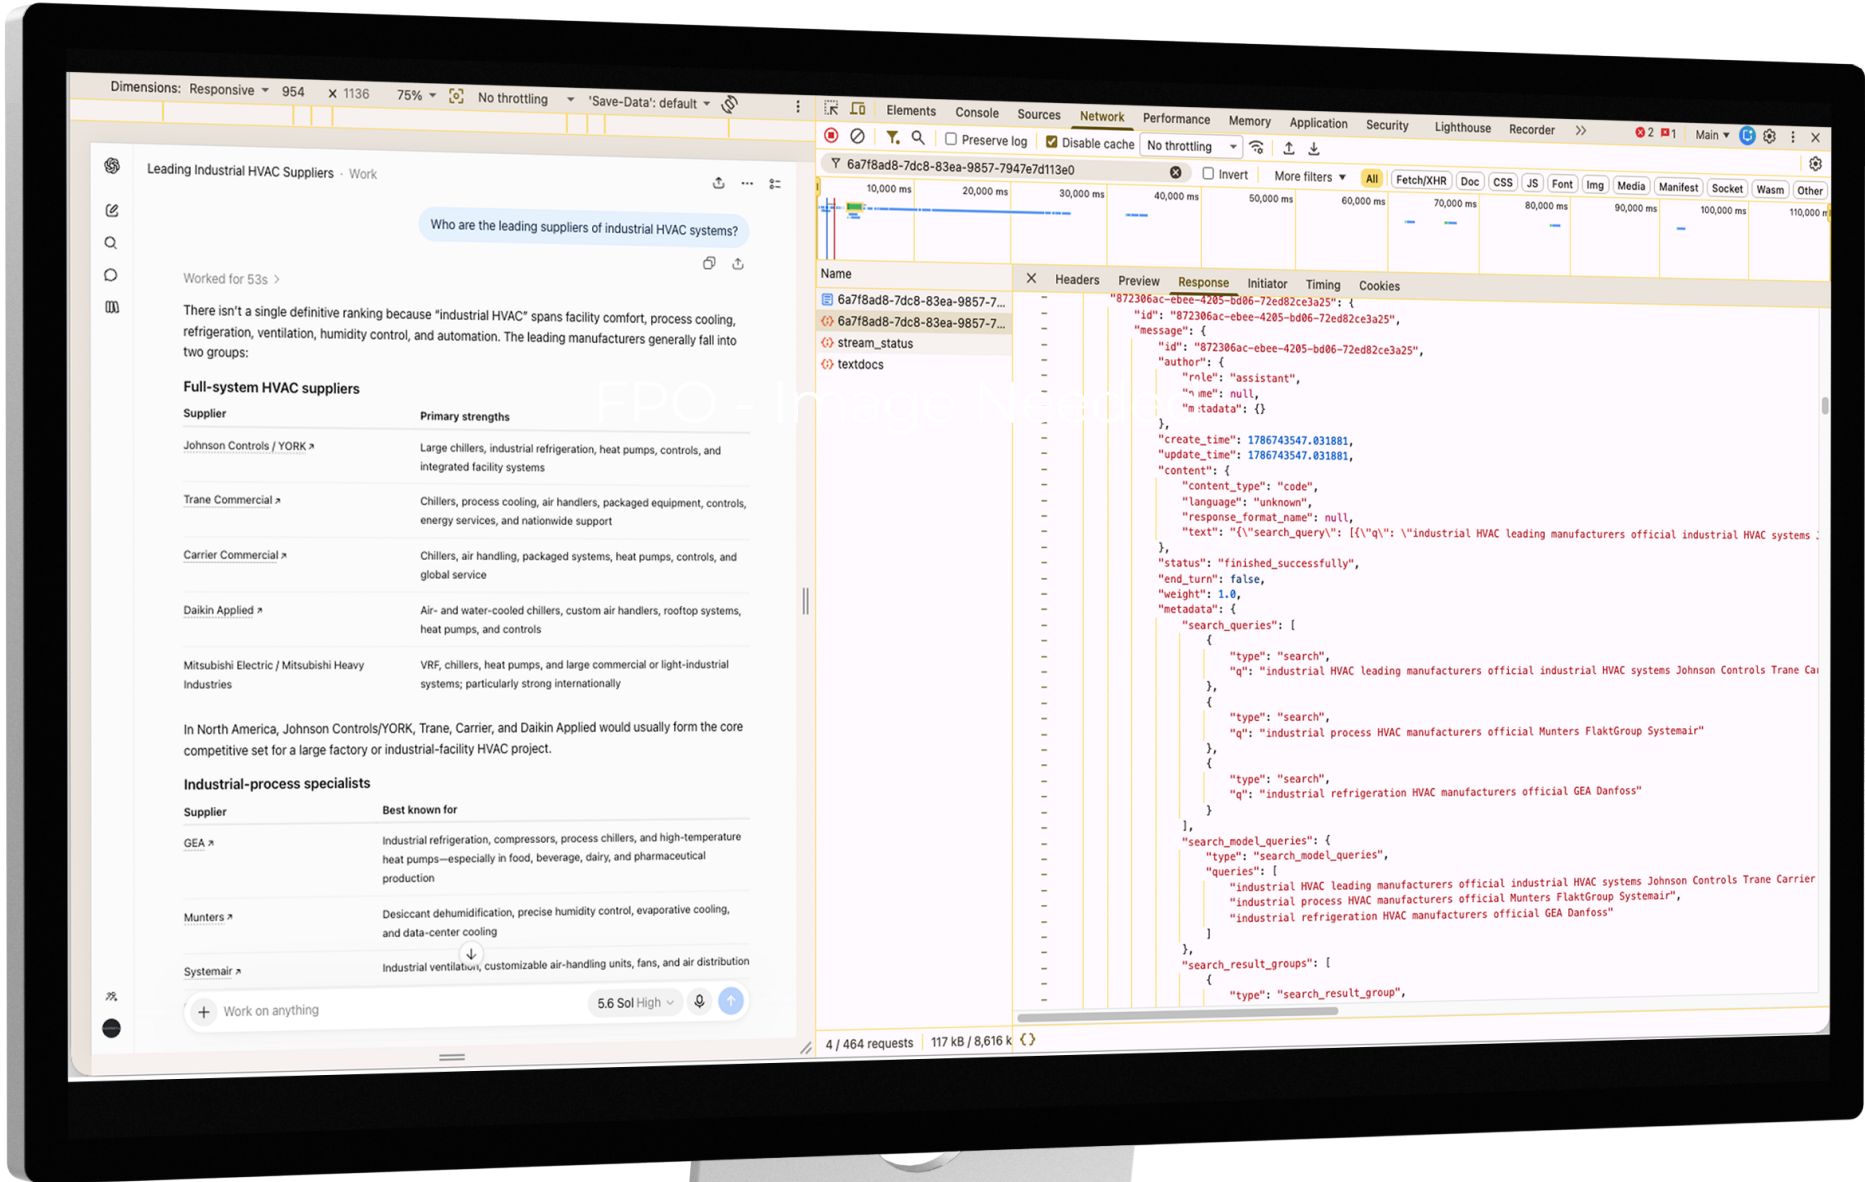
ChatGPT’s goal is to combine the strongest, most consistent information across sources into a clear answer.

As a result, you can appear if you don’t rank #1 in Google — and you might not be included even if you do rank #1. It also explains why third-party sites often show up instead of brand sites.

ChatGPT pulls what forms the clearest answer, not necessarily what ranks the highest.

Let’s look at how this happens. Anyone using the Chrome browser can inspect ChatGPT’s network logs to watch:

- ▶ The actual fan-out searches ChatGPT runs
- ▶ The URLs it pulls back
- ▶ The snippets it extracts
- ▶ The metadata it evaluates



How to watch ChatGPT’s fan-out searches in real time

- 1. In Chrome, navigate to ChatGPT and enter a prompt: “Which HVAC brand is best for my house? I have moderate weather on the East Coast. I want best price and reliability.”
- 2. Copy the Chat ID from your browser URL, which is in a format that looks like this: 6925bce2-6858-832f-aff1-a972d36142bf
- 3. **Right-click** anywhere on the page and click **Inspect**.
- 4. Navigate to the **Network** tab.
- 5. Paste your Chat ID into the filter/search bar (Ex: 6925bce2-6858-832f-aff1-a972d36142bf).
- 6. **Refresh** the page.
- 7. You’ll see rows generating. Look for the request with the **orange icon {;}** and click it.
- 8. Click the **Response** tab.

And just like that, you can see ChatGPT’s full retrieval + reasoning payload.

What You Can Look For

Note: It might be easier to copy/paste that long text into a text editor or Word Doc.

1

Search Queries

First, you can search for the word “queries.” These are the exact fan-out queries ChatGPT triggered from your prompt. This reveals how the model interprets your question. It will look like this below:

```
“search_model_queries”: {
“type”: “search_model_queries”,
“queries”: [
“best HVAC brands for home heating cooling
reliability and price what to look for HVAC brand
reviews”,
“which HVAC brands are considered reliable
for residential HVAC on east coast and how to
choose HVAC brand reliability price”
]
}
```

From that example, you can see ChatGPT fired up two searches with really long descriptive keyword phrases.

2

Snippets + Metadata

Next, search for ‘snippets’ or ‘metadata’ if you want to see what ChatGPT pulled back to analyze. You’ll see:

- ▶ **URLs** – The website page URLs it could potentially cite.
- ▶ **Titles** – The title tags of each of the website pages.
- ▶ **Snippets** – The meta description on the page, or text chunk pulled from the page, if the meta description was blank, long or not useful.
- ▶ **Published Dates** – If a published date could be determined, there’s an epoch timestamp, which is calculated from seconds since 1/1/1970.

In our example below, 1718353440.0 translates to 6/14/24 4:24AM EST. (You can use this [fancy converter](#).)


This is exactly what ChatGPT saw before forming its answer.

```
“metadata”: {
“search_result_groups”: [
{
“type”: “search_result_group”,
“domain”: “www.bobvila.com”,
“entries”: [
{
“type”: “search_result”,
“url”: “https://www.bobvila.com/articles/best-hvac-brands/”,
“title”: “The Best HVAC Brands for Every Budget - Bob Vila”,
“snippet”: “A reliable and efficient HVAC system keeps your home comfortable in all weather. Choose from one of the best HVAC brands with this list of the top picks.”,
“ref_id”: {
“turn_index”: 0,
“ref_type”: “search”,
“ref_index”: 0
},
“pub_date”: 1718353440.0,
“attribution”: “www.bobvila.com”
}
]
}
```

Next, we'll show you the results.

The fan-out query found 30 individual results across 16 unique websites:

- ▶ Bobvila.com
 - ▶ Lifestoryresearch.com
 - ▶ Modernize.com
 - ▶ Pickhvac.com
 - ▶ Meyerair.com
 - ▶ Thefurnaceoutlet.com
 - ▶ Todayshomeowner.com
 - ▶ Heaterguides.com
- ▶ Toocoolair.com
 - ▶ Callmattioni.com
 - ▶ Hookagency.com
 - ▶ Homedepot.com
 - ▶ Hvacblogpost.com
 - ▶ Hvacmind.com
 - ▶ Invoicefly.com
 - ▶ Wikipedia.org



 Why This Matters

This is where visibility begins. If your brand isn't appearing in prompt responses when live searches are used, check whether:

- ▶ It appears in the fan-out queries
- ▶ Your website's metadata is malformed, weak, unclear or missing
- ▶ Competitors dominate multiple queries
- ▶ Your content is being pulled into the retrieval set but not making it into the final answer

It shows that metadata is still critically important and useful in AI-assisted discovery. The same signals that help search engines understand a page — clear titles, descriptions, dates and other page-level details — also help AI systems interpret what content is about as they retrieve, summarize and evaluate information.

This manual process above is useful because it shows you what is happening behind the scenes for a single prompt. But it's not practical if you want to easily understand how often your brand is showing up, which competitors appear more frequently, which sources are being cited, and how those patterns change across many prompts and over time.

Now that we understand how ChatGPT spreads its net and gathers information, the next question becomes: What happens after ChatGPT collects these URLs?

Want to inspect this faster?

Godfrey's [AI Visibility Inspector](#) is a Chrome Extension we built for marketers, strategists and SEO professionals to turn this manual process into a one-click inspection.

Find out exactly which fan-out queries ChatGPT and Claude generate for any prompt and whether your brand does — or doesn't — show up in retrieved results and citations.

Install it for free from the Chrome Web Store.

A Deep Dive into How ChatGPT Scores Content After Retrieval



After collecting the URLs through fan-out searches, ChatGPT performs a two-phase process to determine which brands make it to the final answer:

First, it fetches and reads the pages behind the scenes.

For each URL, ChatGPT:

- Fetches the HTML
- Strips navigation, ads and noise
- Extracts the main body content
- Breaks it into meaningful text chunks

These chunks then individually move into an internal evaluation step.


 Important:

ChatGPT does **not** store these evaluations anywhere. This is real-time scoring, not a historical ranking index like Google. It scores those chunks across 8 broad pillar signals.

These pillars determine:

- What content is trusted
- What information appears reliable
- Which brands feel credible
- What phrasing and framing get elevated

These signals heavily influence what ends up in the final AI answer.



The 8 general categories ChatGPT uses to score content

There’s some gray area here — just like Google’s algorithm, there’s a layer of secret sauce that OpenAI hasn’t published, which we’ve been slowly reverse engineering. But when we analyze retrieval behavior, compare results across prompts and look at what kinds of pages are being consistently cited, a few high-level scoring themes consistently emerge.

1	What it measures: Does this page directly address the user’s question?	⊕ Strong ⊖ Weak Example: A page titled “Best Industrial HVAC Systems for Manufacturing Facilities.” → Direct match to an industrial HVAC query. Example: A general HVAC overview page that never addresses the specific use case.
2	What it measures: Does this source have credibility in this domain?	⊕ Strong ⊖ Weak Signals: ▶ Recognized brand or organization ▶ Strong backlink ecosystem ▶ High editorial standards ▶ Domain known for expertise Example: A well-known manufacturer, major association or established publisher Signals: ▶ Unknown or spam-looking sites ▶ Overly commercial pages with little substance Example: A thin affiliate site with low editorial signals

3

Consensus

What it measures:

Do multiple independent sources say the same thing?

⊕ Strong

- Signals:
- ▶ Repeated facts across many URLs
 - ▶ Industry standards or established definitions
 - ▶ Frequent expert agreement

Example:

Several independent sites agree
Brand X is a leading commercial HVAC supplier

⊖ Weak

- Signals:
- ▶ Claims no other page mentions
 - ▶ Contradictions compared to the broader corpus

Example:

Only one site claims a brand is “#1” with no support elsewhere

Even if your site doesn’t rank strongly, press releases, reviews, industry lists and news coverage often do — and ChatGPT can retrieve and score those pages. That means AI visibility can come from PR, thought leadership, review sites and directories, not just your website. We’ll explore this further in Section 6.

4

Recency (Freshness)

What it measures:

How current is the information?
(Especially in fast-changing categories)

⊕ Strong

- Signals:
- ▶ Recent publish/update date
 - ▶ Content matching current industry standards or regulations

Example:

A 2024 HVAC product specs page updated for new efficiency standards

⊖ Weak

- Signals:
- ▶ Outdated pages
 - ▶ Statistics older than 2–4 years
 - ▶ Old formats or deprecated technologies

Example:

A 2017 article referencing obsolete SEER ratings

5

Diversity of Coverage

What it measures:

Does this source cover the topic holistically?

⊕ Strong

⊖ Weak

Signals:

- ▶ Explains multiple angles (benefits, drawbacks, pricing, features, variations, FAQs)
- ▶ Provides comparisons or context
- ▶ Serves multiple use cases

Signals:

- ▶ Only covers one narrow aspect
- ▶ One-sentence answers
- ▶ No supporting details

6

Sentiment / Framing

What it measures:

How positively or negatively is the brand or topic framed?

⊕ Strong

⊖ Weak

Signals:

- ▶ Evidence-backed positive framing
- ▶ Neutral or balanced tone
- ▶ Industry recognition

Signals:

- ▶ Unsupported hype
- ▶ Negative framing
- ▶ Overly promotional copy

7

Quality / Noise Control

What it measures:

Is the text clean, readable and useful?

⊕ Strong

Signals:

- ▶ Clear structure
- ▶ Strong writing
- ▶ Minimal distractions

⊖ Weak

Signals:

- ▶ Keyword stuffing
- ▶ Disorganized layouts
- ▶ Excessive ads

8

Meta / Structural Alignment

What it measures:

Does the structure help the AI platform understand the page?

⊕ Strong

Signals:

- ▶ Clear title tag
- ▶ Useful meta description
- ▶ Strong H1/H2 hierarchy
- ▶ Geographic or topical alignment

⊖ Weak

Signals:

- ▶ Missing metadata
- ▶ Vague titles
- ▶ Misaligned targeting

Example: How Two Pages Might Score Against the Eight Pillars

Category	Site A (Strong Brand)	Site B (Weaker Competitor)	Why It Wins
Relevance	☆☆☆☆☆	☆☆	Direct match vs. partial coverage
Authority	☆☆☆☆	☆	Recognized manufacturer vs. unknown site
Consensus	☆☆☆☆	☆☆	Broad agreement vs. isolated claims
Recency	☆☆☆☆	☆	Updated vs. outdated
Diversity	☆☆☆☆☆	☆☆	Full guide vs. narrow overview
Sentiment	☆☆☆☆	☆☆☆	Positive framing vs. neutral
Quality	☆☆☆☆☆	☆☆	Structured vs. cluttered
Structure	☆☆☆☆	☆	Clear metadata vs. unclear

Total Weighted Score

- ▶ **Site A: 35/40** → Highly likely to appear in AI answers
- ▶ **Site B: 13/40** → Unlikely to be cited

This is the general high-level decision logic happening behind the scenes, with the model pulling the clearest, most trusted, most complete picture it can find.

 Why This Matters

In many ways, AI visibility isn’t replacing SEO — it’s layering on top of it.

Everything we’ve always cared about still matters:

- ▶ Clear keyword-rich content
 - ▶ Trustworthy authoritative sources
 - ▶ Fresh and accurate information
 - ▶ Consistent messaging
 - ▶ Solid technical structure
- ▶ Relevance is still rewarded
 - ▶ Authority still matters
 - ▶ Consensus carries weight
 - ▶ Freshness influences trust
 - ▶ Off-site mentions (PR, reviews, listicles) reinforce credibility

But AI introduces an additional step in understanding how content is retrieved, evaluated and interpreted by a different model than Google/ Bing’s ranking algorithms.

And while we can’t see every part of that process — nor easily measure these signals at scale — we can see enough through selective inspection to understand the direction:

What’s changing is the mechanism, not the fundamentals. AI doesn’t rank pages, it reasons across them.

And that shift reframes the question for brands: from “Are we ranking?” to “Are we consistently understood, reinforced, and represented across the web?” We now know how AI generally scores content, but none of that matters if the model can’t properly read or interpret your pages to begin with.

The Rules of AI-Readable Content: Structure, Clarity & Segmentability

05

Now that we understand better how ChatGPT scores content, we'll look at the hidden mechanics that determine whether your content actually survives the retrieval process or gets passed over before the scoring even begins.

Why “AI-readable” content matters now

Google and ChatGPT don't read pages the same way. When ChatGPT fetches a URL, it:

- ▶ Strips layout
- ▶ Removes navigation and ads
- ▶ Flattens structure
- ▶ Breaks the text into segments
- ▶ Scores those segments independently

Visual polish doesn't matter. Structure and clarity do.

Here's what ChatGPT looks at:

1

Segmentability: The Invisible Gatekeeper

ChatGPT turns your page into meaning-based chunks — or segments of intent — it can retrieve, score and summarize. If your content isn't clear, you lose.

Each segment should be able to stand on its own. If it only makes sense because of the image above it, the paragraph before it or the sales message around it, it may not survive segmentation well.

- ▶ Messy segments = low scores
- ▶ Clean segments = high comprehension results in higher odds of citation

2

Structure: The AI Version of “Scannability”

Hierarchy is how AI understands the page. A strong structure includes:

- ▶ Clear H1
- ▶ H2s mapped to real subtopics
- ▶ Short, single-idea paragraphs
- ▶ Lists that group information clearly
- ▶ Tables/comparisons (large language models — or LLMs — love these)

Weak structure can be identified by vague headings, mixed paragraphs and fluffy marketing copy. If the hierarchy is unclear, the meaning collapses during segmentation.

3

Clarity: Density of Meaning > Density of Words

AI rewards content that’s explicit, specific, unambiguous and aligned with broader consensus. That means:

- ▶ Define terms early in content
- ▶ Put the answer first
- ▶ Use precise nouns
- ▶ Include data, specs, examples and back claims with evidence
- ▶ Cut the filler (“best-in-class,” etc.)

In other words, clarity is now a technical advantage.

4

Formatting: Signals That Help or Hurt

- ▶ **Help:**
 - ▶ Table
 - ▶ comparison grids
 - ▶ FAQs, bullets
 - ▶ short paragraphs
 - ▶ descriptive headings
 - ▶ clean structure
- ▶ **Hurt:**
 - ▶ Walls of text
 - ▶ Decorative headings
 - ▶ Text-as-image
 - ▶ Messy PDFs
 - ▶ Hidden content
 - ▶ Long intros with no substance

AI is way less forgiving than humans. A person can skim around clutter, but a model may simply fail to extract the right meaning.

5

The AI-Readable Checklist

When reviewing content, ask yourself:

- ▶ Does each section answer one thing?
- ▶ Would the page still make sense with all formatting stripped?
- ▶ Can each segment stand alone?
- ▶ Is the hierarchy clear?
- ▶ Are the key insights near the top?
- ▶ Are headings specific and descriptive?

If yes: AI-ready.

If no: the model may struggle to understand it, retrieve it, or use it confidently in an answer.

 Why This Matters

Foundational SEO practices still matter, including keyword relevance, on-page optimization, metadata, internal linking, crawlability and overall technical health.

AI discovery doesn't replace those fundamentals, but it adds another layer in which:


- ▶ Structure becomes strategy
- ▶ Clarity becomes a visibility advantage
- ▶ Segmentability becomes a gate to visibility
- ▶ Formatting becomes part of AI optimization

Bottom line:


AI isn't judging whether your content is good. It's judging whether it's understandable, extractable and useful in pieces.

And in a zero-click world where a user may get the answer without ever clicking through to a website, understandable content becomes visible content.


Structured



Chaotic



How PR & Off-Site Signals Quietly Shape AI Visibility



Along with content that’s understandable, other factors shape AI visibility as much as your own site: PR, domain authority, brand reputation, third-party mentions and additional off-site signals. It’s where SEO and PR start merging into one unified discovery engine.

So, let’s look at these PR and off-site signals.

PR is no longer just for awareness

PR now directly shapes whether AI platforms surface your brand in answers. For years, PR and SEO lived in adjacent lanes:

- ▶ PR drove credibility, storytelling and reach
- ▶ SEO drove discoverability and traffic

In traditional SEO, PR supports brand awareness and backlinks. But now, in AI-driven ecosystems, PR feeds multiple visibility signals at once. Both SEO and AI platforms still benefit from strong backlink ecosystems — that part hasn’t changed.

What has changed is how broadly PR now influences authority, consensus, trust and understanding inside AI systems.

ChatGPT, Claude, Perplexity and Google AI Overviews evaluate your entire digital reputation

footprint, in which PR plays a major role. Specifically in B2B, it looks at earned media in trade publications and other news sources, searching for content like:

- ▶ Press releases
- ▶ Executive commentary
- ▶ Thought leadership articles
- ▶ Analyst reports
- ▶ Industry listicles

AI systems consistently retrieve these URLs and use them to form understanding, trust and consensus about a brand. If others are talking about you, the AI systems notice, even if your website isn’t in the mix.

PR feeds the authority pillar

Earlier, we noted that authority comes from:

- ▶ Recognized organizations
- ▶ Strong editorial oversight
- ▶ Citations across reputable publications

PR placements check all three boxes.

A press release on Business Wire or PR Newswire?

- ▶ Recognized
- ▶ Editorially vetted
- ▶ Syndicated across multiple reputable domains

A thought leadership article placed in an industry publication?

- ▶ High-trust domain
- ▶ Expertise-driven
- ▶ Often cited by other sources

When ChatGPT decides to gather supplemental info from Google/Bing search results, these are exactly the kinds of top-ranking URLs it pulls in for consideration.

PR helps drive consensus — one signal AI trusts the most

Large language models rely heavily on agreement across independent sources.

If:

- ▶ Your brand appears in a “top vendors” list
- ▶ A trade publication names your product category leader
- ▶ An analyst mentions your capabilities
- ▶ A press release announces a new solution or acquisition
- ▶ A partner references you in a case study

So, even if your own website isn’t ranking well, AI sees your brand as “validated.” This is why PR has become visibility fuel inside AI ecosystems. AI visibility now reflects brand presence across the web, not just on owned domains.

PR helps when your website isn’t ranking — or isn’t SEO crawl-friendly


In many industries (especially industrial B2B), sites:

- ▶ Are updated infrequently
- ▶ Have complex navigation
- ▶ Contain sparse product copy
- ▶ Lack structured markup
- ▶ Don’t appear in mid-tail search queries

Yet even when a website provides limited signals, ChatGPT can still discover your brand through:

- ▶ Press releases
- ▶ Third-party product listings
- ▶ Industry directories
- ▶ Reseller or distributor pages
- ▶ Executive interviews
- ▶ Trade show coverage

This is why brands with strong PR programs can appear in AI answers even when their own SEO is weak.



PR reinforces brand framing across the web

Messaging consistency across coverage reinforces AI’s understanding of your brand’s identity, category and strengths.

If multiple high-quality independent sources say your company is a leading supplier of X category in Y market known for Z capabilities, and those phrases appear in retrieval pages, they will then shape how AI summarizes your brand. AI systems start treating that as fact.

PR = semantic reinforcement

PR + SEO are now one loop

In the old model, PR and SEO often worked in parallel:

- ▶ PR → Earned awareness
- ▶ SEO → Earned traffic

But in reality, there has always been overlap. PR has long supported SEO through backlinks, authority, referral traffic, brand mentions and credibility signals. What’s different in the AI era is how those signals are used.

New Model (AI Era):

- ▶ PR + SEO → Earned understanding, confidence and inclusion

AI systems are not just looking for pages to rank. They’re gathering signals from across the web to understand what your brand does, whether that story is reinforced by credible sources and how confidently they can include you in an answer.

That broader understanding is what influences:

- ▶ Whether your brand is retrieved
- ▶ How your content is scored
- ▶ Inclusion in answers
- ▶ Citation likelihood
- ▶ Brand affinity inside the model

 Why This Matters

PR is data fuel for AI systems. Because in AI-driven discovery:

- ▶ Your brand can appear even when your own site isn’t top ranked
- ▶ Your credibility is shaped by off-site validation (PR coverage)
- ▶ Industry coverage reinforces authority and consensus
- ▶ Backlinks strengthen both SEO and AI presence
- ▶ Off-site messaging can shape how your brand is described

Backlinks still matter enormously. But now they fuel two ecosystems simultaneously:

1. Google/Bing rankings
2. AI retrieval, authority and consensus signals

This is a foundational shift that’s already happening. So, now that we’ve covered how standalone AI models like ChatGPT are surfacing information, let’s shift gears to look at how the industry leader, Google, incorporates AI into search through its AI Overviews.

How Google's AI Overviews Work & What We Can See Behind the Scenes

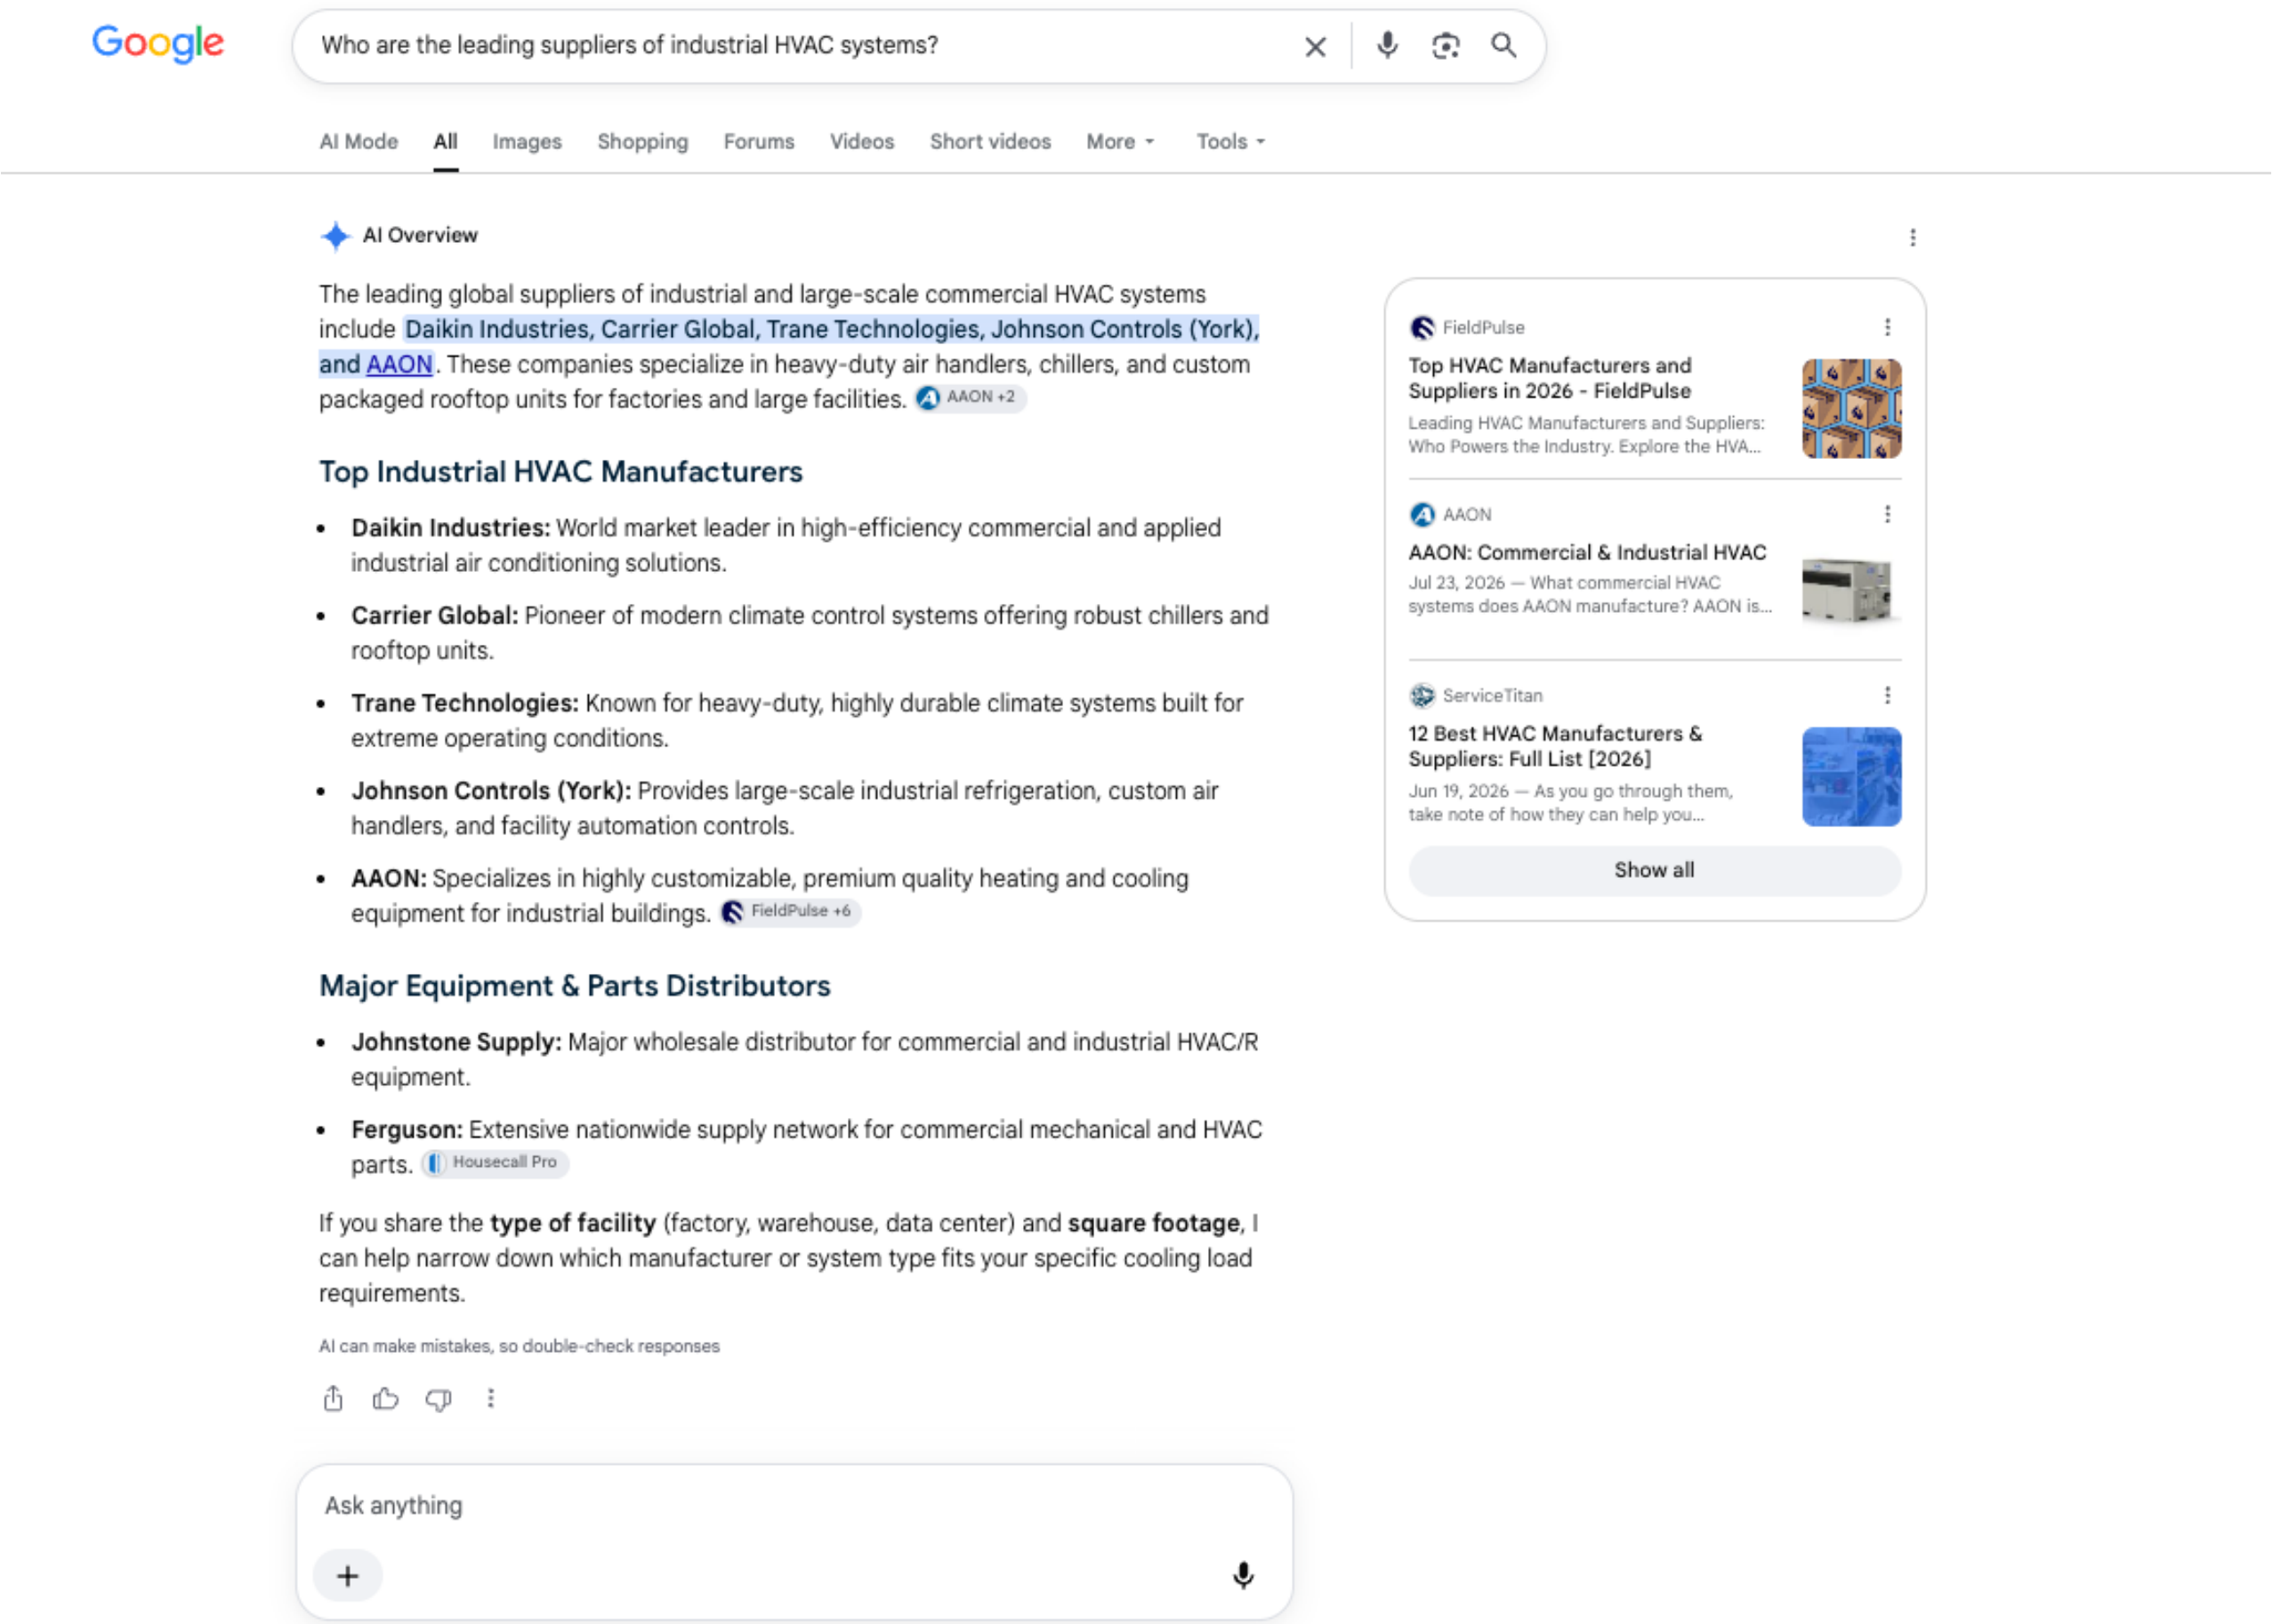


Google remains the leader in online discovery, and AI Overviews (AIOs) are a prime example of how search is evolving.

These summaries pull information from multiple sources, often including top-ranking pages, answer your question on the spot, and increasingly cut out the need to click anything. Depending on the vertical, AIOs now appear on roughly 20–60%, which means they are a meaningful part of the search experience.

But AIOs are only one piece of Google’s broader AI evolution. With AI Mode, Google is moving beyond single answer summaries toward a more conversational search experience, where users can ask follow-up questions, upload files, compare options and conduct deeper research without starting over or ever leaving the search window.

So how do AIOs actually work, and what can we see from Google Search Console (Google’s free platform) that shows how a site performs in search?



What triggers an AI overview

AIOs most commonly appear for:

- ▶ Informational searches
- ▶ How-to queries
- ▶ Product comparisons
- ▶ Troubleshooting questions
- ▶ Longer question-based searches

Google doesn’t generate an AIO for every search. The system evaluates search intent, confidence in the available information and whether an AI-generated answer is likely to improve the user experience.


This is why AIO visibility varies so widely by industry and keyword set. Some verticals see these overviews frequently, while others trigger them far less often.

Why certain brands show up

While Google leans on top-ranking pages for its AIO, it’s not simply about surfacing the pages that rank #1. It can also surface brands through broader authority and relevance signals, including:

- ▶ Strong overall organic visibility
- ▶ Publisher and earned media mentions
- ▶ Industry authority
- ▶ Listicles and comparison content
- ▶ Consistent references across trusted sources

Keep in mind, though, that being mentioned in an AIO doesn’t necessarily mean being cited. Google may reference a brand within a generated answer without directly linking to that company’s website.



How Google decides what to cite

While Google has not publicly disclosed the full weighting model behind AIOs, several factors appear to influence citation selection:

- ▶ **Freshness & relevance related to the search.** Pages that are current and directly aligned to the query have a better chance of being pulled into the answer.
- ▶ **Structured data and clear markup.** Structured data is essentially a set of coded “labels” you add to your site that help Google clearly understand what each part of a page means.
- ▶ **Useful snippets.** Google can pull a single sentence or section from your page if it directly answers the question.
- ▶ **Strong page structure.** Descriptive headings, FAQs, comparison tables, definitions and concise explanations make it easier to understand, extract and cite.

The hidden part: Search console data

One of the biggest challenges with AIOs has been measurement.

Historically, Google Search Console included AIO impressions and clicks, but they haven’t been explicitly labeled and are mixed with regular search activity data. That makes it difficult to know whether changes in impressions, clicks or click-through rates were tied to traditional rankings, AIOs or broader search result changes.

In June 2026, [Google began rolling out](#) dedicated generative AI performance reporting in Search Console, giving site owners more visibility into how their content appears in AI-powered search experiences. Early reporting is still limited, but Google is providing a step toward clearer AI search measurement.

In the meantime, by applying regex filters — smart text filters for your query data — you can identify searches likely to trigger AI Overviews. These filters help isolate “AI-shaped” queries, such as “how to,” comparison searches, or longer, question-based phrases, allowing you to track early AIO visibility for your brand.

Semrush and other third-party SEO platforms are evolving to monitor AIO visibility across tracked keywords, helping connect traditional SEO reporting with emerging AI visibility measurement.

Where this is going: AIO expansion

- ▶ Expect AIOs to keep expanding as Google continues testing new ways to display AI-generated answers across the search experience, including AI Mode.
- ▶ Zero-click searches will keep growing.
- ▶ Google has strong reasons to push this: better user experience + protecting ad revenue.

We’ve just learned more about the ins and outs of AI search systems, but it’s also important to understand how this impacts organic traffic.

The zero-click reality & what the data actually shows

Now that we’ve unpacked how AI systems retrieve, score and surface information, we’ll zoom out to the macro question many marketers are wrestling with:

If more answers now live directly in Google and other AI platforms, then what’s really happening to organic traffic? Between AIOs, featured snippets, calculators and instant answers, it feels like clicks should be disappearing.

But when we look at the data — from Google, third parties and even our own clients’ performance — the story is more nuanced than the headlines suggest.

The zero-click trend is real — but not new

Zero-click searches didn’t just start with AI. Long before AIOs, Google was already answering questions directly in the search engine results pages (SERP) through:

- ▶ Featured snippets
- ▶ Knowledge panels
- ▶ Maps and local packs
- ▶ Calculators and conversions

Third-party data from [Datos](#) and [Pew Research](#) consistently show that roughly two-thirds of Google searches now resolve without a traditional website click.

That sounds alarming — until you look closer at which searches those are.

“How do I make brownies?” vs. real-world B2B search

Some queries are inherently self-contained. If someone searches: “*How do I make brownies?*” Google gives them a recipe:

- ▶ Ingredients
- ▶ Bake time
- ▶ Temperature
- ▶ Instructions

Their job is done. No follow-up action is required. AIOs work extremely well here. But most commercial/industrial B2B searches don’t behave that way.

For example a user might search for:

- ▶ “best industrial hvac supplier for manufacturing”
- ▶ “automatic door systems for hospitals”
- ▶ “precision flow control manufacturers for semiconductor production”
- ▶ “process water treatment solutions for industrial facilities”

Even if Google provides an AI summary:

- ▶ The user still needs details, like specs, manuals, CAD files, etc.
- ▶ They still need to compare options
- ▶ They still need to contact someone

AI can orient the user, but it rarely completes the B2B journey on its own.

These are the kinds of queries our clients live and breathe, and they don’t lend themselves well to single-paragraph, zero-click answers.

It’s critical that your brand still appears in AI summaries to make the consideration set, even though these queries aren’t often fully resolved there. Showing up is what keeps you in the running, even if the buyer has more steps to take before making a decision.

What the data actually shows

A few important data points worth grounding this in:

- ▶ [Semrush research](#) shows that while keywords with AI Overviews do have higher zero-click rates, that rate has not been increasing. In fact, it declined steadily through 2025, reaching ~37.6% by October. This challenges the assumption that AIOs automatically kill clicks.
- ▶ [Datos data](#) reinforces that AI-driven answers are most prevalent in:
 - ▶ Informational queries
 - ▶ Simple how-to searches
 - ▶ Early-stage discovery

They are far less dominant in complex, high-consideration categories.

- ▶ Our own client data mirrors this, as we’ve seen wide variances in how often an AIO appears. Some clients have ~15–20% of tracked keywords triggering AIOs, where others are closer to 50%+.
- ▶ Despite that spread, we haven’t seen broad, systemic declines in organic search traffic tied directly to AIO presence.

AI visibility is not a universal traffic suppressor. That’s why many B2B brands aren’t seeing broad organic declines, even as AI summaries expand.

A key difference: Google vs. ChatGPT


One important distinction that often gets overlooked is that Google still aggressively links out in AIOs.

AIOs frequently:

- ▶ Include multiple inline citations
- ▶ Surface several sources within the answer
- ▶ Encourage deeper exploration with sidebar links (especially on desktop)

This is intentional, because Google is quite good at UX, deeply invested in maintaining user trust, and financially dependent on keeping users inside a healthy search ecosystem. Ad revenue from search represented \$224.5B, roughly 56% of Google’s income in 2025, and they don’t want to lose search market share.

So rather than end journeys, Google wants to guide them more efficiently. That’s very different from ChatGPT’s interface-first, answer-complete design.



Why This Matters

The zero-click narrative is complex. Datos data shows Google still accounts for ~90%+ of global search activity, even as AI platforms expand. Discovery still overwhelmingly starts with Google, and AI is being layered into that experience instead of replacing it.


What is clear:

- ▶ Visibility earlier in the journey matters more.
- ▶ Being understood matters as much as being clicked.
- ▶ Appearing in summaries, lists and comparisons influences where users go next.
- ▶ Traffic hasn’t disappeared — rather, intent has shifted.

Zero click doesn’t mean zero opportunity. It means visibility, credibility and clarity matter earlier in the journey.

And that brings us to a big contributor to what’s dominating AI visibility at the moment, and that’s the “best of” lists.

Why “Best Of” Lists Dominate AI Answers & What To Do About It



With Google's AIO, rich snippets, multi-step reasoning and AI platforms like ChatGPT, Claude and Perplexity generating answers, visibility now happens inside the results themselves, and not just after the click.

So, now we turn to one of the biggest contributors to those AI-generated answers: a content format that zero-click ecosystems rely on more than almost anything else:

“Best Of” lists, rankings, top-10s, vendor roundups, comparison pages: the internet is filled with them. And AI systems love them.

Why? Because as we mentioned earlier, LLMs are obsessed with structure. And listicles are structured by default.

Let’s break it down.

Why LLMs love lists

To a large language model, list-based content is the perfect information format:

- ▶ **Highly structured**
H1 → subheads → bullet points → product summaries These create predictable chunks that LLMs can easily extract.
- ▶ **Easy to analyze**
Lists follow simple, repeated content blocks, which makes them ideal for retrieval.
- ▶ **Keyword-rich**
“Best [category] vendors,” “Top [industry] solutions,” “[Product] vs. [Product]” These phrases match common search terms.
- ▶ **Repeated across multiple domains**
Affiliate sites, review sites, reseller blogs, Reddit, Quora, etc. Repetition = consensus and that’s one of the things LLMs currently value most.

Fit LLMs’ pattern recognition logic

Models aren’t judging editorial quality. They’re recognizing patterns across the web. If “Top 10 ____” pages exist everywhere, the model assumes: “This must be a reliable format for answering the question.” That’s why listicles always show up in Perplexity, ChatGPT, Gemini and Google AIO.

The problem with low-quality lists

While AI loves lists, not all lists are created equal.

- ▶ **Outdated rankings**
Many “best of” lists haven’t been updated in months or years but still rank and get picked up.
- ▶ **Affiliate-heavy content**
A large number of lists exist solely to drive commissions (are paid placements), and don’t reflect actual expertise.
- ▶ **Non-experts shaping visibility**
Your category may be defined online by people with:
 - ▶ No industry experience
 - ▶ No real evaluation criteria
 - ▶ No hands-on product knowledge

Unfortunately, LLMs will still treat these sources as valid if the content structure is strong. This creates a distorted reality: AI repeats what the internet says about you, even if the internet is mostly low-quality listicles.

Google has spent a tremendous amount of time and effort over the last two decades on algorithms to filter spam and sites that ‘game the system’ from claiming top organic rankings. But, unfortunately with AI platforms right now, it seems quantity often matters more than quality.

How to compete ethically

Consider a more strategic approach that gives AI something better, and more authoritative, to retrieve. Here’s a checklist:

- ▶ **Publish “ownable” comparison pages**
Transparent, criteria-based comparisons that explain:
 - ▶ How solutions differ
 - ▶ What specs matter
 - ▶ When each product is the best fit

This establishes your brand as the conclusive voice in your category.
- ▶ **Create transparent, criteria-based lists**
A “Top X Considerations When Choosing ____” guide beats “Best Vendors” and avoids pay-to-play traps.
- ▶ **Make comparison criteria explicit and easy to parse**
Clearly explain how items were evaluated, including the criteria, attributes compared, and why certain options are a better fit for specific use cases.

Use descriptive headings, tables, comparison grids, FAQs and relevant structured data where appropriate, making the comparison easy for both readers and search systems to understand.

It’s one of the cleanest ways to feed trustworthy, machine-readable clarity.

- ▶ **Influence inclusion through outreach**
If authoritative “best of” or comparison lists don’t mention your brand, that’s not the end of the story. Many of these pages are actively maintained and open to updates if you take the right approach.

Thoughtful outreach to site owners, editors and publishers can:

- ▶ Correct omissions
- ▶ Add missing brands, products or categories
- ▶ Update outdated rankings
- ▶ Introduce clearer evaluation criteria

Start by identifying the authoritative sources already shaping the category and showing up in AI answers, then engage with them directly.


- ▶ **Reinforce consensus through credible repetition**
If multiple credible sites repeat your brand/ product’s category name, differentiators and product strengths, AI begins treating those signals as established and trusted. This is how you shift the narrative away from low-quality sources.

How AI uses list content

AI doesn’t just summarize listicles, it “disassembles” them:

- ▶ **Chunk extraction**
LLMs grab discrete product descriptions, attributes or ranking blocks.
- ▶ **Attribute-level comparisons**
If lists say Product A is “fastest” and Product B is “most durable,” AI may mix these attributes into cross-product comparisons.
- ▶ **Indirect visibility via mentions**
Even if your brand isn’t the focus of a question, simply being listed increases inclusion likelihood. In the AI era, being included anywhere increases your chances of being included everywhere.

With a better understanding of how AI looks at your entire digital footprint, we can turn to how AI platforms turn that into a “brand story” — and the challenges that can create.



Message Drift: How AI Platforms Blend Your Brand Story Across the Web



The sheer volume of lists, reviews and scattered sources across the web presents a deeper issue for AI to navigate.

AI blends your entire digital footprint into a single “brand story.” And if your messaging isn’t consistent across the web, that story can drift or become distorted.

Called message drift, this is one of the most important, and least understood, dynamics of AI-era brand building.

1

How AI forms a brand narrative

AI’s job is to gather, evaluate, and synthesize information. When it comes to brand identity, that looks like this:

- ▶ **Retrieval → scoring → synthesis**
Models pull in pages from across the web, score them for authority and then merge them into one logical description.
- ▶ **Semantic reinforcement**
If the same phrase appears across multiple sources (“leading provider of ____”), AI treats it as factual.
- ▶ **Repetition → “truth”**
The more often AI sees a concept or phrase, the stronger that concept becomes in its narrative model. AI isn’t asking: “Is this correct?” It’s asking: “Is this repeated, consistent and reinforced by multiple independent sources?” This is where drift begins.

2

Where drift comes from

Message drift happens when old, inconsistent or third-party content begins outweighing your current messaging.

- ▶ **Old messaging still ranking**
Legacy taglines, old capabilities, retired products are still indexed, still visible, still retrievable.
- ▶ **Inconsistent product descriptions**
Variations across site pages, reseller listings, PDFs, marketplaces, etc.
- ▶ **Distributor or reseller content**
These pages often outrank your own and create conflicting descriptions.
- ▶ **Outdated press releases**
Old positioning signals remain in the ecosystem forever unless updated or displaced.
- ▶ **Internally conflicting language**
If marketing, PR and sales describe you differently, AI ends up blending all of it.

Drift is often cumulative; the result of many small inconsistencies compounded across channels and sources.

3

Examples of drift

When drift happens, AI-generated answers may:

- ▶ **Distort your positioning**
Suddenly you’re a “software company” when you’re actually a service provider.
- ▶ **Put you in the wrong category**
AI may infer a category, sometimes incorrectly, based on old or inconsistent data.
- ▶ **Misrepresent capabilities**
“If they were known for X in 2019, they must still be known for X.”
- ▶ **Mix legacy + current messaging**
One answer may combine your 2021 value proposition, your 2023 product spec, and your 2025 press release boiler plate — all fused into one narrative that may no longer reflect the brand you are today.

4

How to counter drift

The good news: drift can be prevented and, in many cases, reversed.

- ▶ **Create a unified messaging architecture**
A single, consistent framework that governs:
 - ▶ **Positioning**
 - ▶ **Terminology**
 - ▶ **Product naming**
 - ▶ **Category definitions**
- ▶ **Ensure PR + SEO reinforce the same story**
AI relies on both on-site AND off-site references. If your earned media and website stories don't match, drift happens fast.
- ▶ **Repetition across channels**
The more consistent the story, the faster AI models "lock in" the correct interpretation. Correcting drift is less about rewriting everything and more about reinforcing the right signals everywhere, consistently.

Next, we shift from how AI perceives your brand to something even more actionable: *measuring AI visibility and which signals we can realistically track today.*



The Solution

Reinforce the right signals, everywhere, consistently.

How We Can Start Measuring & Optimizing AI Visibility Today



Marketers can measure AI visibility in several ways.

While traditional SEO measurement consists largely of tracking keyword rankings and organic traffic, AI search measurement means tracking how often, accurately and favorably your brand/content appears inside AI-generated answers. This shifts away from quantitative measurements of rankings and traffic to more qualitative measurements like visibility across AI engines and influencing sentiment based on the way these platforms present your brand in its answers.

Let’s dive into some AI visibility measurement tools & methods

Here are some practical ways marketers can diagnose, measure and improve visibility in AI-generated answers, even when users never visit your site.

1 Inspecting fan-out query logs

What this tells you:

How AI systems break a single user question into multiple internal searches before generating an answer.

For example, one purchase-related question posed to AI may trigger multiple background queries (for definitions, comparisons, attributes, pricing, alternatives). Reviewing those fan-out queries helps show how the AI platform interprets the user’s intent and how it pulls content from various sources and synthesizes a response.

When to use it:

- ▶ Your brand doesn’t appear in AI answers for a high-value query
- ▶ A competitor does, and you need to understand why

What to do next:

Use what you learn to adjust your content around the way AI systems are actually searching. That might mean creating clearer comparison content, strengthening metadata, answering related sub-questions or improving the page structures being retrieved.

Want to inspect this faster?

Godfrey’s [AI Visibility Inspector](#) is a Chrome Extension we built for marketers, strategists and SEO professionals to turn this manual process into a one-click inspection.

2

Google search console filters

What this tells you:

Which search queries may be shifting towards AI answers — and where clicks are being replaced by summaries.

Search Console remains important because Google’s generative AI features are still connected to the broader Search ecosystem. While Google’s reporting continues to evolve, marketers can use query filters to isolate patterns that often align with AIO and AI Mode, such as comparison, definition, “best,” “vs,” and question-based searches.

When to use it:

- ▶ Impressions are rising but traffic is flat or declining
- ▶ Click-through rate changes unexpectedly
- ▶ Stakeholders ask why search visibility is not producing the same traffic it used to

What to do next:

Review queries with high rankings and strong impressions, but low clicks. Then, optimize key pages for not only rankings, but for answer inclusion: clear definitions, comparison tables, concise summaries, explicit brand positioning and supporting evidence.

3

Prompt probing frameworks

What this tells you:

When, where and how your brand appears inside AI-generated answers

This approach uses a fixed set of realistic buyer prompts across tools like ChatGPT, Gemini, Claude, Perplexity and Google AI experiences. The goal is to see whether your brand appears, how it is described, which competitors are mentioned and whether the answer reflects your preferred positioning.

When to use it:

- ▶ You need a repeatable way to measure your brand’s AI visibility over time
- ▶ Your internal teams ask, “Do we show up in AI answers?”
- ▶ You want to compare your brand against competitors across various categories
- ▶ You need to track how AI platforms describe your brand, category, strengths or weaknesses

What to do next:

Create a repeatable prompt set tied to real buyer use cases, then run those prompts across priority AI platforms on a regular cadence. Track whether your brand appears, how it’s described, which competitors show up and which sources are cited. Over time, use those patterns to identify content gaps, strengthen PR signals and correct positioning or message drift that may be influencing how AI describes your brand.

This approach is best used for spot-checking perception and competitive framing over time, not exhaustive measurement.

4

SEO tools adapted for AI visibility

What this tells you:

How well your content is structured for AI retrieval, classification, and summarization

Traditional SEO tools are still useful, but the lens changes. Instead of only asking whether a page ranks, marketers should also ask whether the page is clear enough to be understood, useful enough to be summarized and structured enough to be included in an AI-generated answer.

When to use it:

- ▶ Pages rank well in organic search but do not appear in AI summaries
- ▶ AI answers misclassify, oversimplify or omit your offering
- ▶ You need to improve metadata, headings, internal links, content gaps or page clarity
- ▶ You want to make content easier to retrieve, interpret, summarize and cite

What to do next: Use SEO tools to find the inputs that may be limiting AI readability:

- ▶ Unclear titles and descriptions
- ▶ Weak heading structure
- ▶ Crawl or indexability issues
- ▶ Thin sections
- ▶ Missing internal links
- ▶ Outdated content
- ▶ Pages that don't directly answer buyer questions

Then apply an AI visibility lens to improve the content itself. Make the category clearer, define important terms, add concise summaries, answer related questions, strengthen comparison content and reinforce consistent positioning across priority pages.

The goal is to use the data SEO tools already provide to make your content clearer, more accessible and easier for AI systems to understand.

5

Dedicated AI visibility monitoring platforms

What this tells you:

How your brand actually shows up inside AI-generated answers across all AI platforms at scale over time

This is different from reviewing your own website or SEO data. AI visibility monitoring looks at the answer itself: whether your brand appears, how it is described, which competitors show up, which sources are cited, and whether the answer is accurate, favorable, incomplete, or missing you entirely.

And this is where AI visibility measurement gets complicated.

We don't know the full universe of prompts buyers are using. Prompts are longer, more conversational and less predictable than traditional keywords. AI platforms also don't expose clean impression, ranking or click data the way Google Search Console does for traditional search. Even when tools can track AI visibility across a defined set of prompts, the data is still partial. It depends on the prompts selected, the platforms monitored, the timing of the test and how the AI system responds in that moment.

In addition, there isn't one perfect platform that combines traditional SEO performance, Google AIO, ChatGPT, Gemini, Claude, Perplexity, citations, prompt behavior, competitive visibility and optimization recommendations into one complete view. Some emerging tools are starting to solve pieces of the problem, but the space is still evolving.

That means the real value is in knowing which prompts to test, what the results mean, why competitors may be showing up, which sources appear to influence the answer and what to fix first.

When to use it:

- ▶ You want to know whether your brand is showing up in AI-generated answers
- ▶ Manual prompt testing is becoming too limited or inconsistent
- ▶ Leadership wants a clearer view of AI visibility and competitive presence
- ▶ You need to understand how competitors are being described compared with your brand
- ▶ You want to identify practical steps to improve AI visibility over time

What to do next:

Start with a focused AI visibility snapshot. Develop and organize questions, topics and prompts most connected to your business, then evaluate whether your brand appears, how it is described, which competitors show up and which sources seem to influence the answer.

From there, turn the baseline audit into an action plan. That may include strengthening priority website content, improving metadata and page structure, clarifying category language, pursuing PR and third-party mentions, and monitoring results over time to gauge impact.

For many B2B brands, that is the right place to start: not with a massive measurement platform, but with a clear view of where you stand today and what would most improve your visibility next.

What this means for measurement:

AI visibility clearly exists, but measuring it cleanly, consistently, and at scale is still evolving.

For now, these methods give marketers a practical way to diagnose where they are missing, understand what may be influencing AI-generated answers, decide what to fix and track whether visibility changes over time.

Traditional SEO asked:

“Did we rank and did they click?”

GEO asks:

“Did the AI understand us well enough to include us in the answer?”

So, what does that leave us as B2B marketers?

What All of This Means Moving Forward



In this e-book, we’ve peeled back the layers of how AI is reshaping search, from how answers are generated, to how brands show up inside them, to what we can and can’t measure today.

Now it’s time to step back and answer the bigger question: **what does all of this actually mean going forward?**

AI is expanding search, not replacing it

AI platforms are evolving quickly. And buyers are using them. But AI is not replacing the web, Google, SEO or PR. It’s adding another layer on top of them.

Despite the perception that AI platforms like ChatGPT operate independently, they still depend heavily on the traditional web — and on search engines like Google and Bing — for fresh, reliable information. AI-generated answers are shaped by company websites, search results, publisher content, third-party mentions, reviews, forums, news coverage and structured information across the web.

That means the fundamentals still matter. Clear content, strong technical SEO, credible third-party signals, consistent messaging and authoritative sources all continue to shape whether a brand can be found, understood and included.

The takeaway is simple: AI hasn’t eliminated the discovery layer, but it has changed how that discovery layer gets interpreted.

Measurement will improve but remains directional for now

At the same time, the way B2B decision makers use AI is still not fully transparent. We don’t know how often they are using these platforms, what they’re asking or how much those answers influence decisions. We do not have complete data from the AI platforms themselves, but now we have enough visibility to understand the direction things are heading.

We’re now at a point where brands can begin to:

- ▶ See when AI summaries appear
- ▶ Identify queries and prompts that may trigger them
- ▶ Track how brands and competitors are referenced
- ▶ Diagnose why competitors may show up instead

But measurement is still not definitive. AI platforms don’t expose clean impression, ranking or click data the way traditional search tools do. Prompt behavior is harder to predict than keyword behavior. And AI-generated answers can change by platform, model, timing, source availability and the way a question is asked.

The tools will get better. Integrations between AI platforms and marketing platforms are already starting to emerge. But for now, marketers should treat AI visibility data as directional: useful enough to guide strategy, but not precise enough to treat like traditional rank tracking.

Inclusion factors will keep changing

The sources and signals AI platforms rely on will continue to shift as models change, retrieval systems improve, publishers restrict or license access and platforms adjust how they weight different types of content.

That matters because a brand’s AI visibility today is not guaranteed tomorrow. A source that heavily influences AI answers now may become less important later. A competitor that appears frequently today may lose ground if the model, source set or retrieval behavior changes.

This is not a reason to wait. It is a reason to monitor, adapt and keep improving the signals that help AI systems understand what your brand does, where it fits and why it should be trusted.

The practical lesson is that AI visibility cannot be treated as a one-time audit. It needs to become part of an ongoing search, content and PR strategy.

What this ultimately means

- ▶ Search is not disappearing
- ▶ SEO is not dying
- ▶ PR isn’t optional
- ▶ Google is not going away

AI search is still taking shape, but one thing is already clear: it belongs on the B2B marketing agenda.

Not because SEO is going away or Google is being replaced, but because AI is changing how information gets interpreted before a click.

Over time, winning brands will be the ones that:

- ▶ Are consistently understood across the web
- ▶ Reinforce authority and consensus beyond owned channels
- ▶ Optimize for inclusion, not just rankings
- ▶ Monitor AI visibility as an extension of search, not a replacement for it

The work ahead is about building a presence that is clear enough for AI to read, credible enough for AI to trust and consistent enough for AI to repeat.

And for B2B marketers, the opportunity is clear:

Make your brand easier to find, easier to understand and harder to misrepresent — wherever buyers, search engines and AI platforms are looking.

At Godfrey, this is where we can help. As AI search becomes an increasingly important part of how buyers discover, evaluate and compare brands, understanding how your company is represented is becoming a competitive advantage. The brands that actively shape those signals today will be better positioned to stand out as AI continues to influence search and buying decisions.

We can run an AI visibility snapshot to show:

- ▶ Where your brand stands today
- ▶ How competitors are showing up
- ▶ Which content, technical or credibility gaps may be limiting your visibility

From there, we’ll build a practical program to improve how your brand is found, understood and represented across AI-generated answers.



GODFREY >

**TAKE THE NEXT STEP —
GET IN TOUCH WITH US.**

ARE YOU READY?

