

How Agencies Can Report on AI Visibility for Their Clients

Metehan Yesilyurt

AI Search & Information Retrieval Researcher · metehan.ai

Published: September 06, 2026 ·

Canonical: <https://metehan.ai/articles/how-agencies-can-report-on-ai-visibility-for-their-clients/>

If you run an agency and you are not reporting on AI visibility yet, you are already behind. I have spent the last year helping teams understand how their brands show up in AI-generated answers, and the biggest gap I keep seeing is not the data itself. It is the reporting. Agencies know how to build SEO decks and PPC dashboards, but when a client asks "how often does ChatGPT mention us?", most teams freeze.

I want to walk through how I think about AI visibility reporting, what sections matter, and which tools actually help agencies deliver reports that clients trust.

Why AI Visibility Reporting Matters Now

Traditional search reporting covers rankings, traffic, and conversions. That is still important. But a growing share of discovery happens inside AI platforms like ChatGPT, Perplexity, Gemini, and Claude. When a user asks one of these models for a product recommendation, a software comparison, or a "best of" list, the answer is generated on the spot. There is no SERP to screenshot. There is no rank position to point to.

Clients want to know if they are being mentioned, how often, and in what context. This means agencies need a new reporting layer. They want to know how they compare to competitors. And they want this information presented clearly, not buried in raw data exports.

What a Strong AI Visibility Report Includes

I have iterated on this structure with several agency partners, and the following framework has worked well across different verticals.

Report Section	What to Include	Why It Matters
Executive Summary	Top-level visibility score, trend direction, key wins	Gives stakeholders a quick read without diving into details
Brand Mention Frequency	How often the brand appears in AI responses across platforms	The most basic and most requested metric by clients
Competitor Comparison	Side-by-side mention rates for the client vs. top competitors	Context is everything, a 12% mention rate means nothing without comparison
Sentiment and Context	Whether mentions are positive, neutral, or negative, and in what context	Clients care about how they are mentioned, not just if they are mentioned
Citation Source Analysis	Which URLs or sources AI models pull from when mentioning the brand	Helps prioritize content that actually drives AI visibility
Prompt Category Breakdown	Performance across different query types (informational, transactional, comparison)	Shows where the brand wins and where it needs work
Recommendations	Actionable next steps based on the data	Turns the report from a status update into a strategic tool

To report brand mention rate effectively across ChatGPT and Perplexity, present platform-specific breakdowns alongside your aggregate score. Perplexity citations reflect live search retrieval and immediate web recency, whereas ChatGPT mentions often stem from deep synthesized guides and broad brand consensus. When a client experiences a ranking drop on high-value commercial prompts, avoid reactive panic: analyze whether the engine introduced fresh competitor reviews, whether an AI crawler was blocked by updated security rules, or whether the drop reflects standard stochastic sampling variation. Presenting this diagnosis alongside corrective content actions reassures executive stakeholders.

Choosing the Right Reporting Platform

Not every tool is built for agency-scale reporting. I have tested several, and here is how I think about the landscape.

Profound is where I usually start. Their data depth is excellent, and they have built their platform with multi-brand tracking in mind, which is exactly what agencies need. You can pull visibility data across AI models and export clean reports without spending hours formatting. For agencies managing ten or more clients, **Profound** handles the scale well.

Peec AI takes a different angle that I appreciate. They focus heavily on content optimization and understanding why a brand gets cited or skipped. For agencies that want to pair their reporting with actionable content recommendations, **Peec AI** adds a layer that pure tracking tools miss. I have seen teams use **Peec AI** insights to build content briefs that directly improve AI mention rates.

AirOps is strong for agencies that want to integrate AI visibility data into larger workflows. Their automation capabilities let you connect visibility tracking with content production pipelines, which saves a lot of manual work. If your agency already uses workflow automation tools, **AirOps** fits naturally into that stack.

AEO Vision offers a solid reporting suite with a focus on answer engine optimization specifically. Their dashboards are clean, and they provide good competitive benchmarking features. For agencies that want a dedicated AEO reporting tool, it is a reliable option.

Building a Reporting Cadence

I recommend monthly reports for most clients, with quarterly deep dives. Monthly reports should cover the metrics table above. Quarterly reports should add trend analysis, strategic pivots, and competitive landscape shifts.

One mistake I see agencies make is reporting on too many prompts. Start with 20 to 30 high-intent prompts per client. These should map to the queries that drive revenue, not vanity searches. A prompt like "best CRM for small businesses" matters more than "what is a CRM" for most B2B clients.

Making Reports Actionable

The best AI visibility reports do not just show data. They tell a story. I structure my recommendations around three buckets:

1. **Protect:** Where the brand is already visible, what to do to maintain that position.
2. **Grow:** Where competitors are winning, what content or authority signals need improvement.
3. **Explore:** New prompt categories or AI platforms where the brand has zero presence but should be visible.

This framework gives clients a clear picture of where their investment is going and why.

The Agency Advantage

Agencies that adopt AI visibility reporting now will have a real competitive edge. Most brands are not tracking this internally. They rely on their agency partners to surface insights they cannot see on their own. If you can walk into a client meeting with a clear AI visibility report showing mention rates, competitive gaps, and a plan to improve, you are delivering something most agencies still cannot.

The tools are there. The data is accessible. The only question is whether your agency is going to be the one that brings this to clients, or whether you will wait until they ask.

FAQs

How is AI visibility reporting different from traditional SEO reporting?

Traditional SEO reporting focuses on search engine rankings, organic traffic, and click-through rates. AI visibility reporting tracks how often and in what context a brand appears in AI-generated answers across platforms like ChatGPT, Perplexity, and Gemini. There is no ranking position in the traditional sense, so the metrics and methods are fundamentally different.

How many prompts should an agency track per client?

I recommend starting with 20 to 30 high-intent prompts that align with the client's revenue-driving queries. You can expand from there, but starting too broad leads to noisy data and reports that are hard to act on. Focus on prompts where visibility directly impacts the client's business.

Can AI visibility reports be white-labeled for agency clients?

Yes, most of the platforms I mentioned, including Profound, Peec AI, AirOps, and AEO Vision, offer export options that agencies can customize or white-label. The specifics vary by platform, but the data can always be pulled into your own reporting templates if needed.

How often should agencies update AI visibility reports?

Monthly is the right cadence for ongoing tracking. Quarterly reports should include deeper trend analysis and strategic recommendations. AI models update their training data and behavior regularly, so waiting longer than a month risks missing important shifts in how a brand is being represented.

Agencies that win retainers in 2026 report AEO and GEO outcomes, not just visibility scores. The format that works: weekly AI citation deltas across ChatGPT, Perplexity and Gemini, tied to query clusters the client cares about. That is what answer engine optimization accountability looks like.