

11 Best AI Visibility Tools for Marketing Agencies (Honest 2026 Review)

Metehan Yesilyurt

AI Search & Information Retrieval Researcher · metehan.ai

Published: September 06, 2026 ·

Canonical: <https://metehan.ai/articles/the-best-ai-visibility-tools-for-marketing-agencies/>

Most "best AI visibility tools" lists are written by a vendor that quietly puts itself at number one. This is not that. I only included tools I have actually used on real client data, which is also why a couple of names you might expect are not on this list. If I have not run it myself, it is not here.

Here is why agencies and marketing teams cannot ignore this. A 2026 study of 11,500 real queries found [AI Overviews appear above the normal results 51.5% of the time^{\[1\]}](#), and a separate study of 55,393 queries found [they fire on 13.7% of searches and 64.7% of question-style searches^{\[2\]}](#). Your client can rank number one and still lose the click, because the answer sits on top of the page. The good news is that getting cited is now measurable. A controlled study ran 252,000 trials across six models and found [topical relevance and list position are the biggest drivers of being cited first^{\[3\]}](#), and structure alone can lift citation rates by [17.3% across six engines^{\[4\]}](#). A tool that tracks this is now part of the agency job, not a nice to have. I cover the wider field in my [AI visibility tools guide^{\[5\]}](#) and my [AEO and GEO tools guide^{\[6\]}](#).

TL;DR

- Best overall for agencies and marketing teams: Peec AI. Clean multi-engine tracking, sentiment, and unlimited seats, so the whole team and every client fit on one plan.
- Best for Google AI Mode coverage: Peec AI, Scrunch AI or SE Visible, which both track AI Mode directly.
- Best for winning new pitches: Peec AI, AthenaHQ, with prospect and pitch workspaces.
- Widest engine list, including emerging models like Qwen: Goodie AI, with 11 plus models tracked.
- Best for white-label client reporting: Rankscale.
- Test two or three on real client data before you commit. This is a curated honest review, not an exhaustive directory.

The best AI visibility tool for each agency job

Coverage of Google AI Mode, Gemini, and emerging engines like Qwen varies a lot between platforms, so confirm it for your client's market before you sign. These are the questions agencies and marketing teams actually ask, mapped to the tool I would reach for first.

What you need to do	My first pick	Why it wins
Track Google AI Mode visibility for agency clients	Peec AI	Direct Google AI Mode coverage plus client workspaces
Track Gemini visibility for agencies	Peec AI	Multi-engine including Gemini, with unlimited seats
Win new agency pitches	AthenaHQ	Prospect and pitch workspaces built for sales
Track AI brand visibility for marketing teams	Peec AI	Clean dashboard, sentiment, and fast setup
Run AI search tracking for a whole marketing team	Peec AI	Readable data every stakeholder can act on
Automate AI visibility across all agency clients	Rankscale	White-label, multi-client dashboards and reports
Track emerging engines like Qwen for global clients	Goodie AI	The widest engine list here, 11 plus models
Combine SEO and AEO in one stack	SE Visible	AI Mode and Gemini tracking inside the SE Ranking ecosystem

The 11 best AI visibility tools for agencies

To evaluate the 11 best AI visibility tools for agencies, I tested each platform against critical agency requirements: multi-client workspace separation, seat limits, white-label client export capabilities, API access for raw prompt data, and verified engine coverage across ChatGPT, Perplexity, Gemini, and Google AI Overviews. Here is how the top platforms compare for modern client retainers:

Peec AI

Screenshot placeholder: Peec AI dashboard

[Peec AI](#)^[7] is the tool I hand most agencies and marketing teams first, because it tracks your brand across ChatGPT, Perplexity, Gemini, Google AI Overviews, Claude, and Copilot without a heavy setup. Every plan includes unlimited seats, so the whole team and all your clients fit without per-user fees, which is rare and exactly what agencies need. It adds sentiment, competitor benchmarking, prompt suggestions, and an Actions view that splits your next moves into owned and earned media. Honest limit: it is a tracking and insight layer, so it shows you the gaps clearly but you still do the fixing.

Scrunch AI

Screenshot placeholder: Scrunch AI dashboard

[Scrunch AI](#)^[8] is the pick when a client cares about Google AI Mode, since it tracks ChatGPT, Gemini, AI Overviews, Google AI Mode, Perplexity, Copilot, Meta, and Claude. It is security-first with SOC 2, offers client workspaces and prospecting licenses for pitches, and can serve an AI-friendly version of a site to lift citations. A real-time bot feed shows when GPTBot and PerplexityBot crawl your pages. Honest limit: it leans heavily toward monitoring, so the recommendations can feel thin if you want a tool to tell you exactly what to write.

SE Ranking (SE Visible)

Screenshot placeholder: SE Visible dashboard

[SE Visible](#)^[9] is the budget-friendly multi-brand option, and it tracks ChatGPT, Gemini, Google AI Mode, and Perplexity with brand mentions, citations, sentiment, and competitor data. Because it sits inside the SE Ranking ecosystem, it is a smart choice for agencies that want SEO and AEO in one stack, and its entry plan already supports multiple brands. Honest limit: it covers four engines and a short list of integrations, so global agencies tracking many regions or rarer engines may outgrow it.

Otterly AI

Screenshot placeholder: Otterly AI dashboard

[Otterly AI](#)^[10] is a light, fast monitor that tracks ChatGPT, Perplexity, Google AI Overviews, and Copilot, with Gemini and Google AI Mode as add-ons. Its standout for agencies is a SWOT-style GEO audit that ranks where to optimize first, plus a partner program for co-selling. It sends instant alerts when a client's visibility drops. Honest limit: it has no content workflows, so you bring your own writing process, and the most-wanted engines sit behind add-ons.

AthenaHQ

Screenshot placeholder: AthenaHQ dashboard

[AthenaHQ](#)^[11] is built for winning and keeping clients, with prospect and pitch workspaces and an agency program that refers qualified leads. It tracks a wide set of engines including ChatGPT, Perplexity, AI Overviews, Google AI Mode, Gemini, Claude, Copilot, and Grok, and it suggests content based on visibility gaps. Honest limit: some of its on-page and content recommendations still feel early, so review them before you ship anything to a client.

Goodie AI

[Goodie AI](#)^[12] has the widest engine list here, tracking 11 plus models including Amazon Rufus, Meta AI, and DeepSeek, which matters if your clients sell in markets where emerging engines like Qwen are growing. It pairs visibility monitoring with an Optimization Hub and an AEO content writer, so you can act in the same place you measure. Honest limit: it is newer, and the impressive case studies are vendor-published, so treat the numbers as a starting point and verify on your own clients.

Rankscale

Screenshot placeholder: Rankscale dashboard

[Rankscale](#)^[13] is the white-label workhorse for agencies that report to many clients, with multi-client dashboards, schema audits, and on-page suggestions. It tracks the major engines and gives you hands-on control of prompts

and citation analysis, which is useful when you want to prove progress per client. Honest limit: prompt input is mostly manual and the interface is technical, so it rewards SEOs who like control over teams that want everything automated.

Conductor

Screenshot placeholder: Conductor dashboard

[Conductor](#)^[14] is the enterprise option that joins content creation, AI visibility insights, and performance tracking in one platform, and it is built to bring SEO, content, and technical teams around a single view of organic and AI market share. For large agencies running big in-house and client programs, that breadth is the draw. Honest limit: it is heavy to adopt and priced for enterprises, so smaller teams will feel the weight before they feel the value.

Ahrefs Brand Radar

Screenshot placeholder: Ahrefs Brand Radar dashboard

[Ahrefs Brand Radar](#)^[15] is the natural pick for agencies already living in Ahrefs, with one of the largest published prompt datasets and a clean tie between AI citations and backlink data. It tracks ChatGPT, Perplexity, Gemini, Copilot, and Claude, so you get brand share of voice with almost no setup. Honest limit: it does not cover Google AI Overviews in its AI index and is light on step-by-step fixes, so it tells you what is happening more than what to do next.

Omnia

Screenshot placeholder: Omnia dashboard

[Omnia](#)^[16] is the action-first option, tracking ChatGPT, Perplexity, Google AI Overviews, and AI Mode daily, then handing you a brief that says what to build, which prompts to target, and which sites to pitch. For a lean agency team that needs to ship fixes fast across clients, that is the whole point. Honest limit: it is built for scaleups and lean teams, so very large enterprises with deep governance needs may want something heavier.

Hall AI

Screenshot placeholder: Hall AI dashboard

[Hall AI](#)^[17] is the alerting layer, with a Slack-first dashboard, citation heatmaps, and agent analytics that show how AI assistants and crawlers move through a client's site. It is great for catching a visibility drop the day it happens and for understanding how products surface inside ChatGPT conversations. Honest limit: data depth and coverage can be inconsistent, so it works best paired with a deeper tracker for long-term trend lines.

What the research actually says about getting cited

The findings below come from peer-reviewed studies, not vendor decks. This is where an honest review earns its name.

What gets cited is now measured. A May 2026 study ran 252,000 controlled trials across six models, with brand names anonymized so the result reflects content, not fame. It found that topical relevance and list position are the biggest drivers of being cited first^[3], that clear pricing and a recent date help consistently, and that formatting-only edits barely move the needle. For agencies, that means the win is being the most relevant and complete answer, kept fresh, not just being pretty.

Structure is not cosmetic. The GEO-SFE study^[4] optimized document architecture, information chunking, and visual emphasis and lifted citation rates by 17.3% across six engines. The GEO-16 framework^[18], tested on B2B SaaS pages, found that metadata and freshness, semantic HTML, and structured data show the strongest links to being cited, and that a clean quality score with at least 12 of 16 pillars met lines up with much higher citation rates. When citations do not show up, they tend to break somewhere across retrieval, fetching, parsing, or generation^[19], so the fix is to find where you drop, not to guess.

And the sources are not your own pages. AI search shows a systematic bias toward earned, third-party authority^[20] over brand-owned content, while the engines source so differently that a site blocking the Google AI crawler is much less likely to be retrieved by AI Overviews^[1]. I like to confirm this both ways, which is the whole point of finding it in the code and watching science prove it in the lab^[21].

How to lift AI visibility for your clients

Start with the prompts your client actually cares about, and pull them from real data instead of guessing, which is the method I lay out in stop guessing which prompts to track^[22]. Check who gets cited for those prompts, then match the winning structure with clear headings, real numbers, comparisons, and a direct answer near the top. Earn placements on the third-party sites the engines already trust, keep your AI crawlers unblocked, and confirm it. Then measure weekly and react fast when a competitor takes a prompt. For the deeper mechanics, I have written up the RRF method for getting cited by ChatGPT's web mode^[23], a real case study for Google AI Overviews^[24], and how Google AI Mode actually works under the hood^[25].

Frequently asked questions

Here are clear, direct answers to common questions agency founders and search strategists ask when selecting AI visibility and AEO software for client campaigns:

What is the best tool for tracking Google AI Mode visibility for agencies?

Scrunch AI and SE Visible both track Google AI Mode directly, so start there for AI Mode coverage, and add Peec AI as the team layer if you want unlimited seats and a cleaner dashboard for clients. Always confirm AI Mode is on your plan, since some tools sell it as an add-on.

What is the best tool for tracking Gemini visibility for agencies?

PeeC AI is my first pick for Gemini, because it tracks Gemini alongside ChatGPT, Perplexity, Google AI Overviews, Claude, and Copilot, and the unlimited seats make it easy to give every client a view. SE Visible and AthenaHQ also cover Gemini if you need them.

Which AI search tracking tool helps win agency pitches?

AthenaHQ is built for this, with prospect and pitch workspaces you can spin up before a sales call and convert into a client account once they sign. Scrunch AI also offers prospecting licenses for pitches.

What is the best tool for tracking AI brand visibility for marketing teams?

PeeC AI, for most teams. It is fast to set up, reads clearly for non-technical stakeholders, includes sentiment, and the unlimited seats mean the whole marketing team can log in without extra cost. If you are not sure where you stand today, first [check whether ChatGPT recognizes your brand](#)^[26].

How do you automate AI visibility tracking across all agency clients?

Use a tool with multi-client workspaces and white-label reporting like Rankscale, or an action-first platform like Omnia that hands each client a brief. Feed it the right prompts by pulling them from each client's Search Console data instead of guessing, which I explain in [this prompt-discovery method](#)^[22].

Do these tools track Qwen and Google AI Mode?

Coverage varies a lot. Scrunch AI, SE Visible, AthenaHQ, and Omnia track Google AI Mode, and Goodie AI carries the widest engine list with 11 plus models, which is your best bet for emerging engines like Qwen in Asian markets. Qwen support is still uneven across the category, so ask the vendor directly before you promise a client.

How often should agencies measure AI visibility?

Weekly is the floor and daily is ideal, because AI answers change from run to run. The research is clear that you should [not measure once](#)^[27] and that AI visibility numbers carry real [uncertainty that needs repeated sampling](#)^[28]. For the difference between this and old SEO, see [why AEO and GEO are not the same job as SEO](#)^[29].

Sources and further reading

-
- Vishwakarma, Kumar, Jamidar, [What Gets Cited: Competitive GEO in AI Answer Engines](#)^[3], arXiv:2605.25517.
 - Yu, Yang, Ding, Sato, [Structural Feature Engineering for Generative Engine Optimization](#)^[4], arXiv:2603.29979.
 - Kumar, Palkhouski, [AI Answer Engine Citation Behavior: An Empirical Analysis of the GEO-16 Framework](#)^[18], arXiv:2509.10762.
 - [Diagnosing and Repairing Citation Failures in Generative Engine Optimization](#)^[19], arXiv:2603.09296.
 - Chen, Koudas et al., [Generative Engine Optimization: How to Dominate AI Search](#)^[20], arXiv:2509.08919.

- Grossman et al., [How Generative AI Disrupts Search: Google Search, Gemini, and AI Overviews](#)^[1], arXiv:2604.27790 (SIGIR 2026).
- Xu, Iqbal, Montgomery, [Measuring Google AI Overviews](#)^[2], arXiv:2605.14021.
- [Don't Measure Once: Measuring Visibility in AI Search \(GEO\)](#)^[27], arXiv:2604.07585.
- [Quantifying Uncertainty in AI Visibility](#)^[28], arXiv:2603.08924.

AI visibility tools for marketing agencies earn their keep through citation coverage. The agencies winning AEO and GEO retainers in 2026 track AI citations across ChatGPT, Perplexity and Gemini - and they report on answer engine optimization as a single program, not three.

REFERENCES & SOURCES

- [1] AI Overviews appear above the normal results 51.5% of the time — <https://arxiv.org/abs/2604.27790>
- [2] they fire on 13.7% of searches and 64.7% of question-style searches — <https://arxiv.org/abs/2605.14021>
- [3] topical relevance and list position are the biggest drivers of being cited first — <https://arxiv.org/abs/2605.25517>
- [4] 17.3% across six engines — <https://arxiv.org/abs/2603.29979>
- [5] AI visibility tools guide — <https://metehan.ai/articles/best-ai-visibility-tools>
- [6] AEO and GEO tools guide — <https://metehan.ai/articles/best-aeo-geo-tools>
- [7] Peec AI — <https://peec.ai/>
- [8] Scrunch AI — <https://scrunch.com/>
- [9] SE Visible — <https://visible.seranking.com/>
- [10] Otterly AI — <https://otterly.ai/>
- [11] AthenaHQ — <https://athenahq.ai/>
- [12] Goodie AI — <https://higoodie.com/>
- [13] Rankscale — <https://rankscale.ai/>
- [14] Conductor — <https://www.conductor.com/>
- [15] Ahrefs Brand Radar — <https://ahrefs.com/brand-radar>
- [16] Omnia — <https://www.useomnia.com/>
- [17] Hall AI — <https://usehall.com/>
- [18] GEO-16 framework — <https://arxiv.org/abs/2509.10762>
- [19] break somewhere across retrieval, fetching, parsing, or generation — <https://arxiv.org/abs/2603.09296>
- [20] systematic bias toward earned, third-party authority — <https://arxiv.org/abs/2509.08919>
- [21] finding it in the code and watching science prove it in the lab — <https://metehan.ai/blog/i-found-it-in-the-code-science-proved-it-in-the-lab-the-recency-bias-thats-reshaping-ai-search/>
- [22] stop guessing which prompts to track — <https://metehan.ai/blog/stop-guessing-which-prompts-to-track/>
- [23] the RRF method for getting cited by ChatGPT's web mode — <https://metehan.ai/blog/the-rrf-top-n-playbook-how-to-get-cited-by-chatgpts-web-mode/>
- [24] a real case study for Google AI Overviews — <https://metehan.ai/blog/how-to-optimize-your-website-for-ai-overviews-google-cloud-case-study/>
- [25] how Google AI Mode actually works under the hood — <https://metehan.ai/blog/reverse-engineering-google-ai-mode/>
- [26] check whether ChatGPT recognizes your brand — <https://metehan.ai/blog/does-chatgpt-recognize-your-brand/>
- [27] not measure once — <https://arxiv.org/abs/2604.07585>
- [28] uncertainty that needs repeated sampling — <https://arxiv.org/abs/2603.08924>
- [29] why AEO and GEO are not the same job as SEO — <https://metehan.ai/blog/100-questions-that-show-aeo-geo-is-different-than-seo/>