The 6 best AI Search monitoring tools

Masters of Search, presented by Radyant · Published Sep 14, 2026 · Updated Nov 1, 2026

Source: https://www.masters-of-search.com/shortlist/best-ai-search-monitoring-tools

The best AI Search monitoring tools in 2026 are Peec AI, Rankscale, SE Visible, Omnia, AthenaHQ and Scrunch AI. Peec AI ranks first for its source classification and server-log data. It is also the lowest-priced published plan for the full job this guide tracks: 150 daily prompts across ChatGPT, Google AI Overviews and Google AI Mode, with AI-referred traffic reporting and recommendations, at $2,940 a year.

Buyers now ask a model instead of a search engine. They type a question, get three or four names back, and start there. If your brand is not one of the names, you never find out. Nothing shows up in Search Console and nothing shows up in your analytics, because there was no click to lose.

That is the gap these tools fill. You give them the questions your buyers ask. They run those questions against ChatGPT, Gemini, Perplexity, Google AI Overviews and the rest, every day, and tell you who got named, in what order, in what tone, and which pages the model read to get there.

Six tools, ranked for a marketing team or an agency buying this without an enterprise contract.

The shortlist, ranked

#	Tool	Best for	Base	Entry price
1	Peec AI	Marketing teams and agencies that want several engines, citation-level detail and server-log data without moving to an enterprise contract	Berlin, with a New York office since 2026	From $95/mo (Starter, 50 prompts). Agency plans from $245/mo
2	Rankscale	Teams that want the widest engine coverage for the least money, and do not mind a small team behind it	Vienna, Austria	Pro $99/mo (1,200 credits), Pro+ $229 (3,000), Growth $385 (5,500), Enterprise $780 (12,000)
3	SE Visible by SE Ranking	Agencies and marketing teams tracking several brands or clients who want all five engines, several projects and programmatic access included at the entry price	Europe, remote	From $99/mo (Basic, 200 prompts, 3 projects). Core $189 (450 prompts, 5 projects), Plus $355 (1,000 prompts, 10 projects)
4	Omnia	Smaller teams that want to start free, track a focused prompt set, and read the results without training anyone	Madrid	From €79/mo (Growth, 25 prompts). Pro €279 for 100 prompts, Enterprise from €499 for 200+. Prompt packs of 25 on any plan
5	AthenaHQ	Teams that want the tool to act on the findings, and need Meta AI covered	United States	Free Essential tier (300 credits). Starter $295/mo, Enterprise custom
6	Scrunch AI	Brands that want site audits and agent traffic analytics next to the tracking	United States	Core $250/mo (125 prompts). Enterprise custom

Prices verified against each vendor's own pricing page on Oct 16, 2026.

The same workload, priced across every tool

Entry prices are not comparable, because they buy different things. So we priced one identical workload across all of them: 150 prompts, tracked daily against ChatGPT, Google AI Overviews and Google AI Mode, with AI-referred traffic reporting and recommendations, on one seat, for twelve months. That is 13,500 answers a month. It is the loop this guide argues for: track the answers, find out why, and tie them to visits.

Tool	Plan this forces	Per month	12 months
Peec AI	Pro, $245/mo (150 prompts, three engines of your choice, AI Referrals and Actions included)	$245	$2,940
AthenaHQ	Starter, $295/mo (3,600 credits covers about a quarter of this)	$295 plus credit add-ons	From $3,540
Rankscale	Growth, $385/mo (Pro+ caps at 12,000 answers, this needs 13,500)	$385	$4,620
Scrunch AI	Enterprise (Core has no AI Mode and caps at 125 prompts)	Not published	Not published
SE Visible by SE Ranking	No AI-referred traffic reporting or automated recommendations published	Not available	Not available
Omnia	No AI-referred traffic reporting published	Not available	Not available

Three of the six cannot be priced for this job. SE Visible and Omnia do not publish AI-referred traffic reporting, so the table shows them as not available. Scrunch AI's Core plan has no AI Mode and caps at 125 prompts, so the full set is an Enterprise quote. AthenaHQ's Starter covers about a quarter of the answers, so its real price sits above $295. Rankscale's traffic analytics were in beta at the time of writing. Annual billing takes 15% to 20% off most of these. We show monthly, because that is the number vendors advertise.

How this was put together

The buyer we ranked for is a marketing team or an agency that wants daily tracking across several AI engines, bought self-serve, without a procurement cycle. That choice does most of the work in the order below. A tool that covers one or two engines at its entry price ranks lower here than one that covers three or more, however long the vendor has been around.

Everything comes from public sources: each vendor's own pricing page, product documentation and changelog, plus funding announcements and public review profiles. Where a vendor publishes nothing on a point, we say so rather than guess.

Prices are the published list price for a single seat, in the currency the vendor publishes. We use the monthly figure, not the discounted annual one, because the annual discount is easy to find and the monthly number is the one you compare on. Where engines cost extra, we say which ones and roughly what they add, since that is where quoted prices and real invoices usually part company.

We checked the engine lists against each vendor's own documentation in October 2026. This category ships weekly, so treat any coverage claim as a snapshot. Several third-party review posts about these tools are published by competing vendors and were wrong on checkable facts at the time of writing, which is why nothing here is sourced from one.

Entry prices are close to useless for comparing these tools, so we also priced one identical workload across all six and published the arithmetic. The workload covers the whole job this guide describes: tracking the answers, finding out why, and tying them to visits. Where a vendor cannot be priced for that workload from public information, the table says so rather than guessing a number.

The order is our judgement about fit for that buyer. Every entry says where the tool is the stronger pick and where it is not, including the one at the top.

Every entry in detail

1. Peec AI

- **Best for:** Marketing teams and agencies that want several engines, citation-level detail and server-log data without moving to an enterprise contract
- **Base:** Berlin, with a New York office since 2026
- **Entry price:** From $95/mo (Starter, 50 prompts). Agency plans from $245/mo
- **Terms:** Monthly or annual, around 15% off annual. No free tier

Every plan includes three engines of your choice, flat, with no credits to budget, and that is part of the reason it sits first for this buyer. Rankscale reads more surfaces, but it bills a credit per answer, so watching more of them multiplies the bill rather than the coverage. You pick any three of ChatGPT, Google AI Overviews, Google AI Mode, Gemini and Microsoft Copilot, from Starter upwards. The part most teams underrate is what sits underneath the visibility chart: every answer is broken down into the sources the model actually read, classified as editorial, corporate, UGC, reference or your own site, so you can see whether you are losing because nobody writes about you or because the pages that do exist are not being retrieved. Server-log integration is the other unusual piece, with eight ways to connect, so you can watch AI crawlers hit your site rather than infer it. Founded in 2025 and grown quickly since, with more than 3,000 customers and a stated $4M ARR inside the first ten months.

Where it is stronger

- Three engines of your choice at the entry price, with no credits to budget
- Google AI Overviews and Google AI Mode tracked as separate surfaces, which matters because they return different answers
- Source classification and gap analysis, so the data points at what to do rather than only at how you are doing
- Server-log crawl data, which almost nothing else in this price band offers
- Agency plans priced in credits with a published formula, so 25 client projects cost what the arithmetic says
- Fact-checking against facts you supply, which catches models repeating outdated claims about you
- SOC 2 compliant, which is usually the first thing an IT security review asks for

What it does

- Engines, as of 16 October 2026: three of your choice on Starter, Pro and Advanced, including ChatGPT, Google AI Overviews, Google AI Mode, Gemini and Microsoft Copilot. Up to 13 on Enterprise
- Data access: CSV exports and an MCP server on every paid plan. REST API from Pro for brands and on every agency plan. Looker Studio connector from Advanced for brands

- AI-referred traffic: AI Referrals joins measured GA4 sessions rather than estimating them
- Recommendations: Peec Actions produces prioritised actions from the source gaps it finds
- Visibility, position, sentiment and share of voice, tracked daily
- Source and citation analysis, with five source classifications and a gap report
- Crawlability checks across 40+ AI bots from 20+ vendors
- Crawl insights from server logs, via Cloudflare, Vercel, AWS CloudFront, Akamai, a WordPress plugin, webhook or CSV
- Brand perception and fact-checking, which flags answers that contradict facts you supply
- AI shopping and product visibility tracking

Website: https://peec.ai

2. Rankscale

- **Best for:** Teams that want the widest engine coverage for the least money, and do not mind a small team behind it
- **Base:** Vienna, Austria
- **Entry price:** Pro $99/mo (1,200 credits), Pro+ $229 (3,000), Growth $385 (5,500), Enterprise $780 (12,000)
- **Terms:** Monthly or annual, 15% off annual. Seven-day free trial on Pro. Unused credits roll over, to 2x on Pro and Pro+ and 3x on Growth and above. Unlimited seats on every plan

A low entry point and the widest coverage: seven AI interfaces and thirteen models behind them, which is the only split of its kind in this guide. A Vienna company, bootstrapped and built by two founders, and it shows in both directions. Releases land close to weekly, and there is no enterprise apparatus behind it. Citation analysis and page audits are the parts people keep for. The catch is that it sells credits rather than prompts, so the entry tier is cheap and the arithmetic is on you: Pro is 1,200 credits a month, which is 4,800 answers, Pro+ is 12,000, and the REST API starts at Pro+.

Where it is stronger

- The widest coverage in this guide, at seven interfaces and thirteen models
- Interfaces and models reported separately, so you can see whether ChatGPT and the model behind it answer differently
- Citation analysis and page audits in the same tool, so measurement and diagnosis sit together
- Credit rollover to 3x on Growth and above, which few credit-based tools offer
- Fast release cadence, visible in a public changelog

What it does

- Engines included on every paid plan, as of 16 October 2026: seven interfaces, ChatGPT, Gemini, Perplexity, Grok, Copilot, Google AI Overviews and Google AI Mode, plus thirteen models behind them including Claude, DeepSeek and Mistral
- Data access: a Looker Studio connector from Pro, a REST API from Pro+ and an MCP server. CSV bulk import and public dashboard sharing on every plan
- AI-referred traffic: Search Console and GA4 analytics, in beta at the time of writing
- Recommendations: AI Page Audit v2, rebuilt in 2026, reporting against five pillars
- Competitor leaderboards and share of voice

- Sentiment analysis
- Citation analysis showing which sources each engine cited
- Query fanout and a sources box breakdown, showing what the engine searched before it answered
- AI commerce and shopping analysis

Website: https://rankscale.ai

3. SE Visible by SE Ranking

- **Best for:** Agencies and marketing teams tracking several brands or clients who want all five engines, several projects and programmatic access included at the entry price
- **Base:** Europe, remote
- **Entry price:** From $99/mo (Basic, 200 prompts, 3 projects). Core $189 (450 prompts, 5 projects), Plus $355 (1,000 prompts, 10 projects)
- **Terms:** Monthly or annual, 20% off annual. 10-day free trial. Unlimited users on every plan

Most tools in this category started from nothing in 2024. This one comes from the SE Ranking team, who have been shipping search data products since 2013 and now process more than 100 million AI answers a month, and it shows in how much of the job is already built. Five engines are included at the entry price and the list does not change as you move up: ChatGPT, Google AI Overviews, Google AI Mode, Gemini and Perplexity on Basic, Core and Plus alike, with no engine add-ons and no model tiering. The only thing the ladder sells is volume. Alongside the visibility score it shows which prompts trigger a mention, where a competitor takes the question instead, and which domains and URLs the models keep citing, which is the list a PR or outreach team can work from. Every result opens into the answer it came from, with the engine, market and date attached, so a figure in a client readout can be handed to the client. Several brands and client projects run in one account.

Where it is stronger

- Five engines on every plan, so the $99 tier reads the same surfaces as the $355 one
- 200 prompts at the entry price, the most volume at an entry price in this guide
- A REST API and an MCP server on every plan, including the cheapest one
- Every number opens into the unedited answer, with the engine, market and date attached
- Projects, users and research brands are included, not billed per domain or per client workspace

What it does

- Engines included on every plan, as of 16 October 2026: ChatGPT, Google AI Overviews, Google AI Mode, Gemini and Perplexity. Claude announced as coming, Microsoft Copilot not covered
- Data access: a REST API and an MCP server on all plans, read and write, plus CSV export and Looker Studio
- AI-referred traffic: not offered
- Recommendations: not offered as an automated engine
- Visibility, share of voice, average position and sentiment, refreshed daily
- Source analysis showing the domains and URLs each engine cites, grouped by media, blogs and forums
- Prompt and topic view showing which questions surface the brand and how each engine answers them
- Competitor benchmarking against auto-suggested competitors, by topic and by source

- Linked and unlinked brand mentions counted separately
- Markets: the US, UK, Canada, France, Germany, the Netherlands and Spain. Configuration in English, French, German, Dutch and Spanish

Website: https://visible.seranking.com/

4. Omnia

- **Best for:** Smaller teams that want to start free, track a focused prompt set, and read the results without training anyone
- **Base:** Madrid
- **Entry price:** From €79/mo (Growth, 25 prompts). Pro €279 for 100 prompts, Enterprise from €499 for 200+. Prompt packs of 25 on any plan
- **Terms:** Monthly or annual, 20% off annual. 14-day free trial, no card required

The easiest way into this category if you are not sure yet whether you need it. A 14-day trial with no card, then €79 a month for 25 prompts. Founded in Madrid in 2024 by Daniel Espejo, previously at Klarna, and funded with a €3.5M pre-seed in October 2025 led by Visionaries Club. Four engines are on by default and you can turn on up to seven. The thing it does better than its size suggests is surface the actual questions people ask AI about your category, which is useful when you are still working out what to track. Citation tracking runs in any language or region, and insight credits never expire.

Where it is stronger

- A real free trial with no card, which is rare here
- The lowest entry price of the euro-billed tools at €79
- Prompt discovery that tells you what to track, not only how you are doing
- Credits that never expire, so a quiet month is not wasted budget
- Simple enough that a small team can read it without onboarding

What it does

- Engines by default on every plan, as of 16 October 2026: ChatGPT, Perplexity, Google AI Overviews and Google AI Mode. Claude, Copilot and Gemini on request, up to seven in total
- Data access: API access and data exports on paid plans
- AI-referred traffic: not offered
- Recommendations: Omnia surfaces the questions buyers ask about your category
- Citation tracking in any language or region
- Competitive intelligence and share of voice
- Sentiment analysis

Website: https://www.useomnia.com

5. AthenaHQ

- **Best for:** Teams that want the tool to act on the findings, and need Meta AI covered
- **Base:** United States
- **Entry price:** Free Essential tier (300 credits). Starter $295/mo, Enterprise custom
- **Terms:** Monthly or annual, $245/mo billed annually. Free tier available indefinitely

Built by two former Google Search and DeepMind engineers, Andrew Yan and Alan Yao, and backed by

Y Combinator on $2.2M. Eleven models are tracked on Starter, including one most tools skip, Meta AI. The free tier covers five. What separates it from the dashboards is that it tries to do the next step for you: an optimisation agent turns the monitoring data into on-page and off-page recommendations rather than leaving you to interpret a chart. There is a genuine free tier at 300 credits, and then a jump to $295.

Where it is stronger

- Covers Meta AI, which most of this list does not
- A real free tier, so you can see your own data before paying
- Eleven engines on Starter, with no per-engine add-ons
- Recommendations and an agent that acts, rather than a dashboard you interpret
- Founding team with search engineering background at Google and DeepMind

What it does

- Engines, as of 16 October 2026: ChatGPT, Perplexity, Google AI Overviews, Gemini and Copilot on Essential; eleven on Starter, adding Google AI Mode, Claude, Grok, DeepSeek, Meta AI and Mistral
- Data access: API access on Starter is a paid add-on, priced on request
- AI-referred traffic: Google Analytics (GA4) and Search Console integrations
- Recommendations: an advanced content optimisation agent producing on-page and off-page actions
- Source attribution showing which pages each engine referenced
- Sentiment monitoring on brand mentions
- SSO and audit logs on Enterprise

Website: https://athenahq.ai

6. Scrunch AI

- **Best for:** Brands that want site audits and agent traffic analytics next to the tracking
- **Base:** United States
- **Entry price:** Core $250/mo (125 prompts). Enterprise custom
- **Terms:** Seven-day free trial on Core. Five user licenses on Core

Core buys 125 prompts and 5,000 responses a month across ChatGPT, Perplexity, Google AI Overviews and Copilot. Agent traffic analytics and site audits sit alongside the tracking, so it covers more of the job than the cheaper options. Nine engines, Meta AI and Google AI Mode among them, sit on Enterprise.

Where it is stronger

- Agent traffic analytics and site audits in the same product as the tracking
- Five user licenses on the entry plan
- Meta AI covered on Enterprise, which almost nothing else on this list does

What it does

- Engines, as of 16 October 2026: ChatGPT, Perplexity, Google AI Overviews and Copilot on Core; nine on Enterprise, adding Gemini, Google AI Mode, Claude, Meta AI and Grok

- Data access: API access and integrations on Enterprise
- AI-referred traffic: agent traffic analytics
- Recommendations: site audits alongside the tracking, five a month on Core
- 125 unique prompts and 5,000 responses a month on Core
- SSO and a dedicated account team on Enterprise

Website: https://www.scrunchai.com

Engine count is the number that decides your bill

Every tool here publishes a headline price and almost none of them means the same thing by it. Peec AI is $95 and tracks three engines of your choice on 50 prompts. SE Visible is $99 and tracks five on 200. Rankscale is $99 and sells you 1,200 credits, which works out at 4,800 answers a month whichever engines you point them at. Those numbers sit within a few dollars of each other in a spreadsheet and buy very different products.

So before you compare prices, write down the engines you actually need. For most B2B teams that is ChatGPT first, then Google AI Overviews and AI Mode, then Perplexity. Google is worth thinking about carefully, because AI Overviews and AI Mode are different surfaces that return different answers to the same question, and a tool that reports one number for Google is hiding half the picture.

Then price the tools against that list rather than against each other. A cheap tool that charges per engine can end up costing more than an expensive one that includes them, and the gap usually shows up on the second invoice rather than the first.

Measuring is the easy half

Every tool in this guide will tell you that you are mentioned in 30% of answers. Far fewer will tell you why, and that is the number you need, because visibility is an output. The input is which pages the model retrieved before it answered.

This is why source and citation data matters more than the visibility chart. When you look at the URLs behind a prompt where a competitor got named and you did not, the answer is usually one of two things. Either the cited pages are third-party sites you are missing from, which is a PR and review-profile problem, or they are your own pages that the model read and did not quote, which is a content problem. Those need completely different work, and a tool that only reports a percentage cannot tell you which one you have.

A worked example. You sell project management software. You ask ChatGPT for the best project management tool for agencies and you are not named. You open the sources and find four of them are roundup posts on other people's blogs and one is a Reddit thread. None of them mention you. No amount of rewriting your own homepage fixes that, because your homepage was never in the retrieval set. The work is getting into those five pages.

What we left out, and why

We left out the free prompt checkers. They run one question once and show you a snapshot, which is fine for a demo and useless for a trend. AI answers move day to day, and one reading tells you nothing about direction.

We left out the AI visibility tabs inside the big SEO suites. They are the easiest thing to buy, because the invoice and the login already exist, and that is most of what they have going for them. They tend to report Google as one number rather than AI Overviews and AI Mode separately, they price per domain and per seat, and the source-level data is thinner than anything in this guide. If you already pay for one, look at it before you buy. Do not pick your tracker because of it.

We left out the general brand-monitoring platforms that added an AI tab. Social listening and AI answer tracking look similar on a slide and behave differently underneath. If a tool was built to count mentions across the web and now also reads ChatGPT, the AI half is usually shallow, and the source-level data is the first thing missing.

We left out content-optimisation tools that report AI visibility as a side feature. They are often good at what they were built for. They are not measurement tools, and buying one instead of a tracker leaves you optimising against numbers nobody is checking.

We also left out several tools with respectable products and no published pricing at any tier. That is a defensible way to sell software and a bad fit for a guide whose whole premise is comparing what things cost.

Questions buyers actually ask

What should an AI Search monitoring tool cost?

Between $79 and $500 a month for a single brand, depending mostly on how many engines and prompts you need. Published entry points in this guide run from free to $295. A realistic figure for one brand tracking a serious prompt set across five or six engines is $250 to $500 a month. Agencies should price on the credit-based plans instead, because anything billed per domain or per client workspace becomes the largest line item somewhere around the fifth client.

How many prompts do we actually need to track?

Most teams need 50 to 150. Start with the questions your sales team hears, the comparisons buyers make between you and two named competitors, and the category questions where you would want to be one of the three names returned. Fewer than 30 and you are reading noise. More than 200 and you are usually tracking variations of the same question, which costs money and tells you nothing new.

Is Google AI Overviews the same as Google AI Mode?

No, and this trips people up. AI Overviews is the summary that appears above the classic results for some queries. AI Mode is a separate conversational surface you switch into. They draw on different retrieval and regularly return different brands for the same question. If a tool reports a single number for Google, ask which of the two it measured.

Will tracking improve our visibility on its own?

No. These tools measure the problem, they do not fix it. The improvement comes from what you do with the source data: getting into the third-party pages models retrieve, fixing your own pages so they are quotable, and building review profiles that models trust. Buy a tool to know where you stand and to check whether the work is landing. Budget separately for the work.

How long before we see a change in AI answers?

One to two quarters for a shift you can trust. Retrieval indexes refresh on their own schedules, earned coverage takes time to land and longer to be picked up, and review profiles build slowly. Read a set of prompts over weeks rather than one prompt over days. Daily movement in this data is mostly noise, and the tools that chart it hourly are not doing you a favour.

Do these tools work outside English?

Coverage varies more than the marketing suggests. Most track multiple countries and languages on paid tiers, but the engine availability differs by market and some engines are not offered in some countries at all. If you sell outside the US, check the specific country and language list against your markets before you buy, and ask how the tool handles a market where an engine is unavailable.

How much does Peec AI cost in 2026?

Peec AI costs $95 a month for Starter (50 prompts), $245 for Pro (150 prompts) and $495 for Advanced (350 prompts), as checked on 16 October 2026. Enterprise is custom. Every plan includes three engines of your choice and unlimited users. Three engines is the plan limit, so tracking more than three costs extra. Some roundups list a 25-prompt plan at €89. Peec AI has no such plan.

Is Peec AI only a monitoring tool?

No. Peec AI tracks visibility, position, sentiment and share of voice, and it also classifies the sources behind each answer, lists prioritised actions in Peec Actions and shows AI crawler visits from server logs. Some roundups call it monitoring-only, which leaves those parts out.