The 6 best AI visibility tools for agencies

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

Source: https://www.masters-of-search.com/shortlist/best-ai-visibility-tools-for-agencies

Every roundup of these tools prices them for one brand. Agencies do not buy them for one brand. They buy them for ten, or forty, and the pricing model is what decides whether that works or quietly eats the retainer.

This is where the two halves of the category separate. Some of these tools charge per domain and per seat, which is fine for a brand and brutal for an agency. Others sell a pool of usage you spread across client projects, which is the model that scales. The feature lists look nearly identical either way.

Six tools, ranked and priced for an agency running ten client brands.

The shortlist, ranked

#	Tool	Best for	Base	Entry price
1	Peec AI	Agencies that want per-client projects, predictable credit maths and six engines without per-domain billing	Berlin, with a New York office since 2026	Agency plans from $245/mo (Essential, 3 client projects)
2	Rankscale	Agencies that want the widest engine coverage per dollar and can live with a two-person vendor	Europe, remote	Pro $99/mo, Growth $385, Enterprise $780. An Essentials tier is listed at $20 with no published credit allowance
3	AthenaHQ	Agencies that want the tool to draft the optimisation work, and need Meta AI or Amazon Rufus for retail clients	United States	Free Essential tier (300 credits). Starter $295/mo, Enterprise custom
4	Keyword.com	Agencies who want to prove the problem to a client before selling them a programme	United States	From $3/mo for the rank tracker. AI add-on from $24.50/mo for 50 credits
5	Semrush AI Visibility Toolkit	Agencies already running every client on Semrush who want one login and one invoice	Boston, with offices worldwide	$99/mo per domain. Semrush One, bundling SEO and AI visibility, from $199/mo
6	Gauge	Agencies whose bottleneck is producing the content, not finding the gaps	San Francisco	From $99/mo (Starter, ChatGPT only). Growth $599

Prices verified against each vendor's own pricing page on Sep 14, 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: ten client brands, 35 prompts each, tracked daily against six engines, for twelve months. That is 350 prompts and 63,000 answers a month. Thirty-five is the middle of the 25 to 50 range we recommend per client further down, and pricing the bottom of that range flatters whichever tool has the cheapest credit tier.

Tool	Plan this forces	Per month	12 months
Peec AI	Agency Scale, $795/mo (65,000 credits, 25 client projects)	$795	$9,540
Semrush AI Visibility Toolkit	$99 per domain plus a prompt block, ten domains	From $1,590, for four engines	From $19,080
Gauge	Growth, $599/mo (no multi-client plan published)	Not published	Not published
Rankscale	Above Enterprise (top published tier caps at 48,000 answers)	Not published	Not published
AthenaHQ	Enterprise (Starter includes 3,600 of the 63,000 credits needed)	Not published	Not published
Keyword.com	Credit add-on, sold in blocks of 50	Not published at this scale	Not published

The agency number that matters is not this one, it is what the eleventh client costs. On Peec AI Scale it costs nothing, because the plan includes 25 client projects and 65,000 credits against the 63,000 this uses. On Semrush it costs another $159 a month, every month, forever. On Rankscale it costs a renegotiation, because a book this size already runs past the top published tier. Rankscale is genuinely cheaper at a lighter load, and at 25 prompts per client it comes in $15 under Peec AI, which is worth knowing if your clients really are that light. Annual billing takes 15% to 20% off most of these.

How this was put together

The buyer is an agency running AI Search tracking for ten client brands, wanting daily data, client-ready reporting, and a bill that does not scale linearly with the client count. Client work is the constraint that matters here, so a tool sold per domain ranks below one sold as a pool of usage, whatever its feature list says.

Everything comes from public sources: each vendor's own pricing page, its documentation, and the agency customers it names publicly. Where a vendor publishes no agency plan, we say so rather than inferring one from its brand pricing.

We priced one identical agency workload across all six and published the arithmetic below. Three of them cannot be priced for it from public information. That is the finding rather than a gap in the research, and it tells you something real about who the product was built for.

The order is our judgement about fit for that agency. Every entry says where the tool is the stronger pick and where it is not, including the one at the top. Where a tool beats the top entry on a lighter workload, the table note says so plainly.

Every entry in detail

1. Peec AI

- **Best for:** Agencies that want per-client projects, predictable credit maths and six engines without per-domain billing
- **Base:** Berlin, with a New York office since 2026
- **Entry price:** Agency plans from $245/mo (Essential, 3 client projects)

- **Terms:** Essential $245, Growth $495, Scale $795, Comprehensive custom. Around 15% off annual

The only tool here with a published plan line built for agencies rather than adapted from brand pricing, and the one with the agency customers to show for it. Seer Interactive, Peak Ace, Mindshare, Kinesso, Eskimoz, Previsible and Omniscient are all named publicly. The model is a pool of credits you spread across client projects, with a formula anyone can check: one prompt, against one model, for one day, costs one credit. That means you can work out next quarter's bill on the back of an envelope before you sign a client, which is the thing per-domain pricing takes away from you. Six engines are included on every tier. Scale is $795 for 65,000 credits and 25 client projects, so a ten-client agency has room to double without a new contract.

Where it is stronger

- The only published agency plan line here, with client projects as a named unit
- Credit maths you can do before signing a client, rather than after
- Scale includes 25 client projects at ten-client prices, so growth does not trigger a renegotiation
- Six engines included, with no per-domain or per-seat multiplier
- Named agency customers at holding-company scale, which no other tool on this list can show
- Looker Studio and API export, so client reporting can live in the deck you already send
- SOC 2 compliant, which is usually the first thing an IT security review asks for

What it does

- Engines included on every plan, as of 14 September 2026: ChatGPT, Google AI Overviews, Google AI Mode, Perplexity, Gemini and Microsoft Copilot
- Data access: CSV exports, REST API and an MCP server on every agency plan. Looker Studio connector from Growth
- 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
- Agency plans with per-client projects: 3 on Essential, 10 on Growth, 25 on Scale
- A published credit formula: one prompt, one model, one day, one credit
- Source and citation analysis with five source classifications, plus a gap report
- Crawl insights from server logs, via Cloudflare, Vercel, AWS CloudFront, Akamai, a WordPress plugin, webhook or CSV
- Brand perception and fact-checking against facts you supply

Website: https://peec.ai

2. Rankscale

- **Best for:** Agencies that want the widest engine coverage per dollar and can live with a two-person vendor
- **Base:** Europe, remote
- **Entry price:** Pro $99/mo, Growth $385, Enterprise $780. An Essentials tier is listed at $20 with no published credit allowance
- **Terms:** Monthly or annual, 15% off annual. Unused credits roll over to a per-tier cap

The best engine coverage per dollar here, and the cheapest option for a genuinely light client book. Enterprise is $780 for 48,000 answers and 100 brand dashboards, which is plenty of dashboards and

not enough answers: at 35 prompts per client across six engines, ten clients need 63,000, so a normal book runs past the top published tier and into a quote. More than 17 engines are covered, against six included elsewhere. Rankscale is bootstrapped and built by two founders, Mathias Ptacek and Patrick Schmid, which suits agencies whose clients buy without a formal vendor review.

Where it is stronger

- Cheapest here if your clients are light, at $780 a month for ten clients on 25 prompts each
- 17+ engines, which is roughly triple what most of this list includes
- 100 brand dashboards on Enterprise, so client count is not the binding constraint
- Credit rollover, so a quiet month funds a busy one
- Per-tier credit maths published in full, including the consumption rate

What it does

- Engines included on every plan, as of 14 September 2026: ChatGPT, Perplexity, Google AI Mode, Google AI Overviews, Gemini, Claude, Copilot, Grok, DeepSeek and Mistral, with 17+ supported in total
- Data access: REST API, Looker Studio connector and an MCP server
- AI-referred traffic: Search Console and GA4 analytics, in beta at the time of writing
- Recommendations: GEO site audits with 94+ technical checkpoints
- Brand dashboards: 10 on Pro, 50 on Growth, 100 on Enterprise
- Credit rollover to a published per-tier cap
- Citation analysis showing which sources each engine cited
- Scheduling from hourly to monthly, across all global regions

Website: https://rankscale.ai

3. AthenaHQ

- **Best for:** Agencies that want the tool to draft the optimisation work, and need Meta AI or Amazon Rufus for retail clients
- **Base:** United States
- **Entry price:** Free Essential tier (300 credits). Starter $295/mo, Enterprise custom
- **Terms:** Monthly or annual, 17% off annual. API access on Starter is a paid add-on

The pick if your bottleneck is producing recommendations rather than producing charts. An optimisation agent turns the monitoring data into on-page and off-page actions, which for an agency is the difference between a report a junior writes up and one that arrives half-written. Coverage is broad and uniform, with every model available on every tier including Meta AI and Amazon Rufus, which matters for retail clients. Built by two former Google Search and DeepMind engineers, Andrew Yan and Alan Yao, on $2.2M and a Y Combinator backing. Usage is credit-based at one credit per AI response, so ten clients sits above the Starter allowance.

Where it is stronger

- A free tier you can use to pilot a client before putting them on a plan
- Meta AI and Amazon Rufus covered, which most of this list does not offer
- Every engine on every tier, with no per-engine add-ons to explain to a client
- An agent that drafts the work, not just a dashboard that describes it

- Founding team with real search engineering background

What it does

- Engines included on every tier, as of 14 September 2026: ChatGPT, Perplexity, Google AI Overviews, Google AI Mode, Gemini, Claude, Copilot and Grok, plus DeepSeek, Meta AI and Mistral
- Data access: API access on Starter is a paid add-on, priced on request
- AI-referred traffic: not offered
- Recommendations: an advanced content optimisation agent producing on-page and off-page actions
- A free Essential tier at 300 credits, usable as a client pilot
- The Athena Citation Engine, on Starter and above
- Multi-country support on Starter and above
- SSO and audit logs on Enterprise

Website: https://athenahq.ai

4. Keyword.com

- **Best for:** Agencies who want to prove the problem to a client before selling them a programme
- **Base:** United States
- **Entry price:** From $3/mo for the rank tracker. AI add-on from $24.50/mo for 50 credits
- **Terms:** Monthly or annual. Free plan for 10 keywords

The cheapest way to put real data in front of a prospect. Three dollars for the core rank tracker, a free plan at ten keywords, then fifty AI credits for $24.50. For a pitch, that is enough to show a client they are missing from answers their competitors win, which is usually the entire argument you need to make.

Where it is stronger

- The cheapest genuine entry point in this category, by a wide margin
- A free plan, so a prospect can be shown their own data at no cost
- A mature daily rank tracker underneath, not a standalone AI widget
- Useful as a new-business tool rather than a delivery tool

What it does

- Engines included, as of 14 September 2026: ChatGPT, Gemini, Perplexity and Google AI Overviews, as a credit-based add-on
- Data access: exports and API on the core rank tracker
- AI-referred traffic: not offered
- Recommendations: not offered, this reports position
- A mature daily rank tracker underneath the AI module
- A free plan at 10 keywords, and $3/mo for 50

Website: https://www.keyword.com

5. Semrush AI Visibility Toolkit

- **Best for:** Agencies already running every client on Semrush who want one login and one invoice
- **Base:** Boston, with offices worldwide
- **Entry price:** $99/mo per domain. Semrush One, bundling SEO and AI visibility, from $199/mo
- **Terms:** Per domain and per seat. Extra prompts and extra seats are billed separately

If your whole agency already lives in Semrush, the argument here is operational rather than technical, and it is a real argument. One login, one invoice, and the AI data sits next to the keyword and backlink data your team already reads. Pricing is per domain, so ten clients comes to $1,590 a month across four engines.

Where it is stronger

- No new vendor, no new login, no new invoice and no new training
- AI data next to keyword, backlink and technical data, read by one analyst
- The largest support and research organisation of anything here
- Client reporting already wired into whatever Semrush reporting you run today

What it does

- Engines included, as of 14 September 2026: ChatGPT, Google AI, Gemini and Perplexity, with Google reported as one combined surface
- Data access: exports and reporting within the Semrush platform
- AI-referred traffic: available through the wider Semrush analytics suite rather than the AI toolkit
- Recommendations: a site audit scoring AI readiness
- Prompt tracking with daily AI rankings, 25 prompts included per domain
- Sits alongside the classic SEO toolkits

Website: https://www.semrush.com

6. Gauge

- **Best for:** Agencies whose bottleneck is producing the content, not finding the gaps
- **Base:** San Francisco
- **Entry price:** From $99/mo (Starter, ChatGPT only). Growth $599
- **Terms:** Monthly or annual. Y Combinator backed, founded 2024

Founded in 2024 in San Francisco and backed by Y Combinator, with a small team and an unusually direct product. Growth is $599 for 600 tracked prompts a day across eight engines, which is a large allowance at that price. The differentiator is a content engine that drafts and publishes articles straight from the visibility gaps it finds, so the output is a draft rather than a chart.

Where it is stronger

- 600 prompts a day on Growth, a large allowance for the money
- Eight engines on the Growth plan
- A content engine that drafts from the gaps, which suits a team short on writers
- The agent and the content engine are included rather than priced as add-ons

What it does

- Engines included, as of 14 September 2026: ChatGPT only on Starter; ChatGPT, Google AI Overviews, AI Mode, Gemini, Perplexity, Copilot, Claude and Grok on Growth
- Data access: exports and an Ask Gauge agent, included rather than an add-on
- AI-referred traffic: not offered
- Recommendations: a content engine that drafts and publishes articles from the visibility gaps it finds
- 600 tracked prompts a day on Growth, which is a large allowance for the price
- Citation analysis across a large corpus of AI answers

Website: https://www.withgauge.com

Per domain or per pool, and why it decides everything

There are only two pricing models in this category and the difference matters more than any feature.

Per domain means you pay a fixed amount for each brand you track. Semrush works this way. The cost rises in a straight line with your client count, which means every new client has a fixed software cost attached before you have done any work. At $99 a domain, a forty-client agency is paying $3,960 a month.

Per pool means you buy a quantity of usage and spread it over as many clients as you like. Peec AI, Rankscale and AthenaHQ all work this way. The cost rises with how much you actually track, not how many logos you have, so a client on a small prompt set costs you very little. It also means you can put a prospect in the tool during a pitch without adding a line to the bill.

For an agency the pool model wins almost every time. The exception is when a single client is large enough to fund its own subscription, which happens, and at that point you should be billing the software through to them anyway.

What your clients will actually ask for

Three things, in our experience, and only one of them is on any vendor's feature page.

First, they want to know why, not what. A chart showing 22% visibility generates one question, which is what to do about it. The tools that answer it are the ones that show you the URLs the model read before it answered, because that tells you whether the gap is a PR problem or a content problem. Those are different budgets and different teams, and guessing wrong costs a quarter.

Second, they want it in the reporting they already get. Nobody wants another login. Looker Studio export, an API, or a CSV that survives contact with your deck template matters more than the dashboard's own design, because your client will only ever see your version of it.

Third, and this one arrives late, they want to know whether it is working. That means a number that moves over quarters rather than days, and a record of what you changed and when. Daily volatility in this data is mostly noise. If you report it weekly you will spend your QBRs explaining variance instead of progress.

What we left out, and why

We left out tools with no published pricing at any tier. An agency has to quote a client before it buys,

and a tool you cannot price is a tool you cannot scope.

We left out the social listening platforms that added an AI tab. They count mentions across the web, which is a different job, and the source-level data an agency needs to act is usually the first thing missing.

We left out white-label resellers of other people's tracking. Several exist, most are a wrapper around one of the engines already on this list, and the margin you gain is paid for in a support chain you do not control.

We also left out the free prompt checkers. They are useful in a pitch and useless in a retainer, because one reading of one question tells you nothing about direction.

Questions buyers actually ask

What should an agency budget for AI visibility tracking per client?

Between $60 and $100 per client per month at ten clients, on a pool-based plan, for a normal 25-prompt set across six engines. On per-domain pricing it is $99 per client before extras, and it does not come down as you grow. Most agencies bill this through to the client as a line item rather than absorbing it, which is easier if the plan lets you show per-project usage.

Can we white-label these tools for client reporting?

Mostly not in the dashboard itself. What you actually get is export: Looker Studio connectors, APIs and CSV, which you then put in your own template. Treat white-labelling as a reporting question rather than a platform question, and check the export path before you buy, because that is the part your client will see every month.

How many prompts should we track per client?

Twenty-five to fifty for most clients. Start with the comparisons their buyers make between them and two named competitors, the category questions where they would want to be one of three names, and the questions their sales team hears every week. Below fifteen you are reading noise. Above a hundred you are usually tracking rephrasings of the same question and paying for it.

Do we need a different tool for enterprise clients?

Sometimes, and it is a compliance question rather than a feature one. An enterprise client's security review will ask about SOC 2, SSO and data residency, and that is where most of this list runs out. Peec AI is SOC 2 compliant, and AthenaHQ publishes SSO and audit logs on its Enterprise plan. None of the others here publishes a full compliance posture, and Rankscale says plainly that it has none. So if a client's security review gates the purchase, start with Peec AI or AthenaHQ, and check before you put any name in a proposal.

Will tracking improve our clients' visibility on its own?

No. These tools measure the problem. The improvement comes from the work the source data points at: getting clients into the third-party pages models retrieve, making their own pages quotable, and building review profiles models trust. Sell the tool as the measurement layer under a retainer, not as the retainer.

How do we prove this is working to a client?

Pick the prompt set at the start of the engagement and do not change it, because changing it changes the number without changing reality. Report quarterly on visibility and average position across that fixed set, alongside what you shipped in the period. Add the source data showing which pages started getting cited. That last part is the one clients find persuasive, because it connects a number they cannot feel to work they can see.