

AI-powered search visibility is quickly emerging as a crucial KPI for enterprise SEO teams, especially in B2B SaaS and multi-location brand environments. Traditional SEO metrics barely scratch the surface when it comes to measuring an organization's presence in AI-driven search ecosystems — be it ChatGPT, Gemini, Perplexity, or newer entrants like Claude and Google AI's own Overviews and Modes. Building a vendor scorecard tailored to AI visibility tools is essential to make informed decisions that go beyond marketing hype.

In this post, I'll share my 12 years of experience in enterprise SEO to help you create a practical, detailed, and sanity-checked scorecard. We'll cover key themes like prompt-level tracking at scale, multi-LLM coverage, citation intelligence, and — most importantly — how to verify pricing without getting stuck behind endless sales demo loops.

Why AI Search Visibility Is a New Enterprise KPI

AI visibility must become a core enterprise KPI because:

- **Search is evolving:** Users increasingly ask AI assistants and LLM-driven engines rather than typing keywords into Google classic search.
- **Traditional tools fall short:** Legacy SEO platforms track indexed pages and backlinks but typically don't surface prompt-level rankings or semantic AI citation data.
- **Multi-LLM ecosystems:** Businesses must measure visibility across different LLMs like OpenAI's ChatGPT, Google's Gemini, Microsoft's Copilot, and others — no single vendor fully covers this today.

Making AI search visibility a standardized KPI ensures you focus on impact where your brand is actually discovered in AI queries. But to do that, you need to create a vendor scorecard to evaluate prospective tools rigorously.

Core Components of a Vendor Scorecard for AI Visibility

Your scorecard should have three primary pillars: **Pricing Verification**, **Feature Checklist**, and **Performance & Coverage**. I'll break down each below.

1. Pricing Verification: Don't Trust "Unlimited" at Face Value

Many vendors tempt enterprise buyers with "unlimited seats," "unlimited exports," or "unlimited prompt tracking." As someone who sits through countless demos, I always *sanity-check* these claims by asking for:

- The exact export limits per user and per month
- Prompt tracking caps, e.g., max number of prompts tracked simultaneously or in a reporting period
- Any hard limits hidden behind "custom enterprise" plans

Without this clarity, you risk budget surprises or hitting critical bottlenecks during launches.

Example: Take *Peec AI* — their pricing includes:

Starter: €89/month, Pro: €199/month, and Enterprise pricing is custom. Knowing the starter and pro tiers lets you frame your expectations, but you must always ask what the "custom" tier entails. Do Enterprise plans truly remove prompt tracking caps? Are seats actually unlimited or capped at 50? Do exports require an additional fee? Get exact numbers before moving forward.

2. Feature Checklist: The Must-Haves for AI Visibility Tools

AI visibility is a new, moving target. Here are key features you want to vet:

1. **Prompt-Level Tracking at Scale:** Can the tool track thousands of prompts in real time? Does it segment by prompt type and user intent?
2. **Multi-LLM Coverage:** Check support for major large language models including:
 - OpenAI ChatGPT
 - Google Gemini
 - Perplexity AI
 - Anthropic Claude
 - Google AI Overviews/Mode (not just traditional search results)
 - Microsoft Copilot (when possible)
3. **Citation/Source Attribution & Intelligence:** Since AI responses heavily rely on cited sources, your tool must analyze the quality, authority, and freshness of citations embedded within AI-generated answers.
4. **Reporting & Dashboards:** Does the platform offer customizable dashboards that marry AI visibility with traditional SEO data?
5. **API Access & Data Export Limits:** For enterprises building custom workflows, API access is essential. Export limits should be clearly stated.

If a tool only tracks Google AI Overviews and claims “AI visibility,” it’s a red flag. AI search is truly multi-LLM — the scorecard keeps you honest.

3. Performance & Coverage

Enterprise-grade AI visibility tools need to demonstrate real-world performance:

- **Frequency of Data Updates:** Are prompt performance metrics updated daily, weekly, or in real time?
- **Geographic & Language Support:** For global companies, coverage in multiple languages and regions is critical.
- **Data Integrity & Sample Size:** Is the data sample statistically significant? Does the platform disclose confidence intervals or margins of error?

Building Your Enterprise AI Visibility Vendor Scorecard: Step-By-Step

Here’s a simple framework you can adopt and tailor:

1. Define Your Goals & Must-Haves

Work with your enterprise stakeholders — SEO, product, analytics — to prioritize which AI models and regions matter most. Identify your minimum feature set and pricing budget.

2. Research & List Vendors

Compile a list of vendors offering AI visibility tools. Don’t just rely on marketing sites — check forums, communities, and your network for recent feedback.

3. Sanity-Check Pricing & Limits

Request detailed pricing breakdowns. Ask each vendor to specify export caps, prompt tracking limits, API call quotas, and seat maximums. Document responses side by side in a table.

4. **Feature Matrix**


Create a matrix listing all features you care about. Mark which vendor supports which feature fully, partially, or not at all. Probe for demo videos or actual dashboard screenshots.

5. **Test Data Quality**

If possible, run free trials or pilots to validate real data freshness, coverage, and prompt-level insights.

6. **Score & Recommend**

Weight each category (price, features, performance) by your priorities and calculate a total vendor score. Share your findings transparently.



Sample Pricing & Features Table

Vendor Pricing (Starter/Pro/Enterprise)			Prompt-Level Tracking	Multi-LLM Coverage	Citation Intelligence	Export Limits	Seats
Peec AI	Starter €89/mo	Pro €199/mo	Enterprise Custom	Yes (scalable)	ChatGPT, Google AI Overviews, Perplexity	Basic citation quality scores	Starter: 1,000 exports/mo Pro: 5,000 exports/mo (approx.) Starter: 3 seats Pro: 10 seats
Vendor X	Starter \$99/mo	Pro \$299/mo					



Enterprise Custom Yes ChatGPT, Gemini, Claude, Copilot Advanced citation source freshness and authority
Unknown (needs sales confirmation) Unlimited claimed (needs verification)

Note: Always demand explicit caps and confirm them in contract terms. "Unlimited" is almost never truly unlimited.

Common Vendor Red Flags to Watch For

- **Marketing buzzwords without substance:** Phrases like "next-gen AI visibility" or "revolutionary insights" without concrete feature demos.
- **Opaque pricing:** Vendors requiring "contact sales" without a clear pricing framework.
- **Limited AI model coverage:** Only Google AI Overviews with no support for other major LLMs.
- **Export and API limits buried:** Key data caps hidden in fine print or verbal only.

Final Thoughts: Show Me the Prompts

In an era where "AI visibility" can easily become empty jargon, your enterprise vendor scorecard must act as your sanity check. Always ask vendors to **show you the prompts** — literally the queries and AI responses their tool tracks — to judge relevance and scale. Demand full disclosure on pricing caps and feature support, especially for multi-LLM coverage and citation intelligence.

By building a thoughtful, data-backed fingerlakes1.com scorecard that combines pricing verification, a clear feature checklist, and real-world performance validation, your team will avoid costly missteps — and ensure your SEO efforts keep pace with AI's rapid evolution.

Need help building your vendor scorecard? Reach out with your scenario, and I can help you establish a no-nonsense, enterprise-grade evaluation framework.