

In today's fast-evolving digital landscape, businesses no longer only compete through classic SEO tactics. The rise of AI-powered search and Large Language Models (LLMs) has ushered in a new era: **AI search visibility**. As an experienced B2B SaaS analyst and former enterprise martech buyer, I've seen firsthand the importance of asking the right questions when evaluating complex AI visibility platforms.

If you're evaluating solutions like *Peec AI*, which starts at **€89/month (Starter)**, moves up to **€199/month (Pro)**, and offers custom pricing at the enterprise level, your due diligence must be rigorous to avoid costly pitfalls. This blog post covers the critical questions you should ask during any AI visibility demo, across several essential themes:

- AI search visibility vs. classic SEO metrics
- Prompt-level measurement and tracking
- Multi-LLM coverage and assistant benchmarking
- Share-of-voice, sentiment, and citation tracking
- Refresh frequency, export options, and scalability concerns

Understanding AI Search Visibility vs. Classic SEO

Classic SEO metrics focus largely on keyword rankings, backlinks, and organic traffic derived from traditional search engines such as Google and Bing. AI search visibility, however, demands a new measurement paradigm. The question to ask is:

1. How Does the Platform Define and Measure AI Search Visibility?

Do they track your brand presence in AI-driven search environments such as ChatGPT, Bing AI, or other LLM-powered assistants? Without a clear, replicable methodology, claims of "AI visibility" can be marketing fluff.

For example, Peec AI offers integrated AI search tracking that includes prompt testing and conversation modeling—but what precisely are the metrics? <https://dailyiowan.com/2026/02/09/5-best-enterprise-ai-visibility-monitoring-tools-2026-ranking/> Can they show you measurable presence based on actual AI output data, or is it extrapolated from standard SEO rankings?

2. What is the *Refresh Frequency* of AI Search Data?

Many demos hype "real-time" AI search data—yet treat "real-time" as once every 24 or 48 hours, which can be insufficient in fast-changing environments. Ask about exact refresh intervals and how that impacts data accuracy. For Peec AI, understanding their data update cadence is crucial to ensure actionable insights.

Prompt-Level Measurement and Tracking

The heart of AI visibility lies in prompts—the inputs given to AI models that generate outputs. Classic keyword tracking is no longer enough; you must drill down to prompt success and usage patterns.

3. Does the Tool Offer Prompt-Level Analytics and Tracking?

Probe how prompts are tracked:

- Can the platform measure prompt engagement and success rates (e.g., whether the AI's output matches desired outcomes)?
- Are prompt variations tested automatically to optimize responses?
- Is prompt attribution transparent, so you know which prompt generated which output?

Some platforms give generic "AI engagement" numbers; you want granular insights. Peec AI's demo should show prompt tracking dashboards and how they handle prompt versioning.

4. How Do They Handle Metrics Like AI Output Quality and Relevance?

These terms abound in marketing, but ask for measurable definitions—precision, recall, sentiment alignment—or at least clear scoring criteria. Vague "AI governance" buzzwords without defined thresholds or examples are a red flag.

Multi-LLM Coverage and Assistant Benchmarking

The AI landscape is fragmented. Multiple LLMs (e.g., GPT-4, Claude, Bard) exist, and users access AI assistants with differing models and interfaces.

5. What Level of LLM Coverage Does the Platform Provide?

Some platforms focus solely on a single LLM, which limits visibility. Confirm if the AI visibility platform covers all major models and emerging ones, including their various API versions and user interfaces.

During your demo, ask specifically:

- Is GPT-4 covered? Earlier GPT models? Other providers like Anthropic or Google's Bard?
- How do they handle API updates and model version changes?
- Do they benchmark model performance for your prompts across those LLMs? If so, how detailed is this comparison?

Peec AI markets multi-LLM capabilities, but you want proof. For enterprise scale, it's important to know whether coverage scales with your increasing footprint.

6. Can You Benchmark Assistant Performance on Your Keywords or Prompts?

Knowing general model capabilities is insufficient. You need benchmarking tailored to your brand or industry keywords to understand share-of-voice among assistants.

- Does the platform let you compare the same prompt across models to see which produces better outputs (accuracy, sentiment, citation integrity)?
- Are these benchmarks updated regularly to reflect model improvements?

Share-of-Voice, Sentiment, and Citation Tracking

Beyond raw visibility, grasping the quality and context of AI mentions is critical. This includes tracking sentiment and sources that AI assistants cite in responses.



7. How is *Share-of-Voice* Calculated in AI Contexts?

Unlike classic SEO share-of-voice—which is a relatively defined metric—AI share-of-voice requires analyzing how often your brand or keywords appear in AI outputs versus competitors.

Ask for:

- Clear methodologies and formulas for share-of-voice calculation specific to AI responses
- Examples and benchmarks for your competitive set
- Historical trends so you see shifting AI visibility over time

8. Does the Platform Track Sentiment Around Your Brand or Keywords in AI Outputs?

AI responses can carry positive, negative, or neutral sentiment that may impact brand perception. Check:

- Which sentiment analysis models are used? Are they custom-tuned for your industry?
- Granularity of sentiment tracking—at prompt, query, or brand mention level?
- Whether sentiment metrics update on each data refresh

9. Are Citations and Source Attribution Monitored? How Are They Tracked?

An increasingly important aspect of AI visibility is whether AI outputs cite your content as a source. This impacts trust and SEO indirectly. Ask:

- Can the platform parse and track citation references in AI responses?
- Are citation patterns broken down by source, assistant, and prompt?
- How are citations weighted in relation to visibility metrics?

Critical Operational Questions: Refresh Frequency, Export Options, and What Breaks at Scale?

It's easy to get dazzled by feature lists and dashboards. But focus on operational realities:

10. What Exactly is the Refresh Frequency of Data? Is It Truly “Real-Time”?

Ask for specifics on how often each data type is refreshed—search results, prompts, sentiment, citations. For large enterprises, a 24-hour delay might be too slow. Confirm how refresh frequency varies across tiers (e.g., Starter vs Enterprise).

11. What Export Options Exist? Are Raw Data, Reports, and APIs Available?

Access control and export capabilities are often buried in demo decks but are make-or-break for workflows. Your questions:

- Can you export raw prompt-level data and AI output logs?
- Are exports available in standard formats like CSV, JSON, or via API?
- Is there granular control over who in your organization can access or export data?

12. What Breaks at Scale? Are There Usage or Query Limits?

Many platforms have limits on prompts analyzed, queries per month, or LLM interactions. For Peec AI, review pricing footnotes for any tiered caps on:

Pricing Tier Monthly Price (EUR) Potential Query or Prompt Limits Starter €89 Check platform for prompt/query caps Pro €199 Higher limits; confirm exact numbers Enterprise Custom Scalable limits with SLAs

Ask what happens when you exceed limits: throttling, additional fees, or service degradation.

Summary Checklist: What to Ask in Your AI Visibility Demo

1. **Definition & Methodology:** How is AI search visibility measured? What’s actually tracked?
2. **Refresh Frequency:** What is the data update schedule for AI search, prompts, sentiment, and citations?
3. **Prompt Analytics:** Can you track prompt-level performance, variations, and outputs?
4. **AI Output Quality:** How are output relevancy and quality measured objectively?
5. **LLM Coverage:** Which models and assistants are covered? How do they handle API/version changes?
6. **Assistant Benchmarking:** Can you compare prompt performance across models?
7. **Share-of-Voice:** How is it calculated and benchmarked?
8. **Sentiment Tracking:** Are sentiment metrics clear, granular, and regularly updated?
9. **Citation Tracking:** Does the platform monitor where your brand/content is cited in AI responses?
10. **Export & Access Controls:** What export options and user permissions exist?
11. **Scalability & Limits:** What limits exist on usage, and what breaks at scale?

Final Thoughts

Evaluating an AI visibility platform requires more than admiring fancy dashboards or on-paper feature sets. You must verify measurable, transparent metrics with clear definitions, adequate refresh rates, and realistic scale considerations. Peec AI’s pricing and positioning are competitive, but the true test lies in hands-on demos asking these probing questions.

By focusing on what is concretely measurable—prompt tracking, multi-LLM benchmarks, citation monitoring, refresh cadence, and export controls—you’ll be equipped to select the platform that genuinely supports your AI visibility strategy in a growing and fragmented AI landscape.

