



CiteHawk

The framework, published in full

— THE BUYER'S STANDARD

The AI Visibility Audit.

What a real audit of your brand's AI visibility should actually contain. We publish the complete structure, the same one our platform runs, because a framework you can check beats a promise you can't. Use it to run your own audit, or to grade anyone selling you one.

An audit that skips any of these six layers is measuring part of the problem and billing for all of it.

The answers: are you in them?

Everything starts with the only question a buyer's AI ever answers: which brands get named. A real audit asks the actual buying questions of your category across all 8 major AI models (ChatGPT, Claude, Gemini, Perplexity, DeepSeek, Grok, Copilot and Google AI Overviews), asked without your brand in the prompt, multiple times, with every answer stored verbatim.

The outputs that matter: **presence** (which questions name you at all), **position** (first mention and top-three placement, not just "mentioned"), and **per-model splits**, because leading ChatGPT while absent from Gemini is the normal case, not the exception. An audit that reports one blended number has averaged away the finding.

Red flag when buying an audit: screenshots of a single chat session. AI answers vary run to run; a single sample is an anecdote. Ask how many runs per question, and whether the raw answers are included.

The competitors: who wins the answers you lose?

Every question you lose, somebody wins. The audit should name them, per question, with their position, and then aggregate: who wins most often, on which AI models, and with what gap. This section usually reorders a marketing team's mental competitor set within one meeting, because the brands AI recommends against you are frequently not the brands you benchmark against.

The facts: what does AI get wrong about you?

Ask the AI models direct questions about your brand: what you do, what you cost, who you serve, where you operate. Then verify each claim against reality. AI models repeat stale pricing, dead features and wrong categories with total confidence, and every buyer who asks hears the same wrong answer.

This layer matters because it is the cheapest fix in the entire program. Wrong facts trace to sources: an outdated directory profile, a stale about page, an old review. Correct the source and the answers follow. An audit that skips fact-checking has skipped the fastest win it could have sold you.

The sources: where do the answers come from?

AI answers cite a recognizable set of sources: category directories, review platforms, comparison articles, community threads, publications. The audit must extract every source the collected answers cite, then split them into **sources that mention you** and **sources that carry your competitors but not you**.

That second list, ranked by how often each source appears in answers, is the actionable core of the whole audit. It converts "improve your AI visibility" from a vibe into a work order: these profiles, these publications, this order.

"An audit without a ranked source-gap list is a diagnosis without a prescription."

The technical read: can AI systems use your site?

When AI models retrieve pages mid-answer, your site either helps or hinders. The technical layer checks, at minimum:

Access: robots.txt exists and does not accidentally block the AI crawlers you want (AI model fetchers and AI-search bots are different animals from training crawlers; blocking policy should distinguish them).

Discovery: a valid sitemap, and an llms.txt that actually describes the site rather than existing as an empty gesture.

Liftability: page-level scoring of how extractable your key pages are: direct answers under clear headings, structured comparisons, plain statements of what you are and what you cost.

Alignment: whether the pages you have actually match the questions buyers ask. A beautiful page answering a question nobody asks scores zero here.

The entity check: does the machine world know you exist?

AI models classify brands partly through the entity layer of the web: whether your brand resolves cleanly to a known entity, whether references to you are consistent and unambiguous, whether your name collides with something else. The audit should check how recognizable your brand is as an entity and flag ambiguity, because a model that cannot confidently resolve who you are cannot confidently recommend you.

What the deliverable must look like

Whoever runs the audit, hold the deliverable to four tests:

- 1 • Receipts attached.** Every claim traces to a stored answer you can open. No screenshots-as-proof, no "our data shows" without the data.
- 2 • Per-model, never blended-only.** One score is a headline; eight columns are a strategy.
- 3 • Ranked actions, not observations.** Facts to fix, sources to win, pages to restructure, in impact order.
- 4 • Repeatable.** An audit is a snapshot. The answers move as models retrain and retrieval refreshes, so the same measurement must be re-runnable on a rhythm: weekly measurement is the honest cadence, with monthly review.

This framework is deliberately public. If a vendor cannot match it, you have learned something. If they can, this document just made their pitch shorter.

Run all six layers on your brand, free.

CiteHawk's first report runs this exact audit across all 8 AI models: presence, competitors, brand facts, source gaps, the technical read and the entity check, with every answer stored as evidence. No card required, ready within the hour.

Start your free audit → citehawk.com/signup

© 2026 ScaleUp Teams Pty Ltd · CiteHawk, the AI brand visibility monitor · Method: citehawk.com/leaderboards/methodology · Reproduction with attribution welcome.