

AI Visibility Audit Kit

The exact prompts we use to test a brand's visibility across ChatGPT, Perplexity, and Google's AI Overviews - plus a scoring sheet and a worksheet for what to fix first.

Before you start

- Use clean sessions: no chat history, no custom instructions, logged-out where possible. Assistants personalise; you want the answer a stranger gets.
- Run each prompt in at least three: ChatGPT, Perplexity, and Google (AI Overviews / AI Mode). Claude and Copilot if your buyers skew technical or corporate.
- Record answers verbatim, same day. Screenshots beat memory.
- Budget an afternoon for the first run. Quarterly re-runs take an hour.

Part 1 - The prompt pack

Swap the brackets for your category, market, and buyer.

Category prompts - do you exist where buying starts:

1. "What are the best [category] providers in [market]?"
2. "I'm a [buyer role] at a [company type] dealing with [problem]. Who should I look at?"
3. "Compare the leading [category] options for a [company size] business."
4. "Who are the most credible [category] specialists for [industry]?"
5. "Recommended [category] agencies/tools for a budget of [range]?"

Comparison and objection prompts - what happens mid-decision:

1. "[Your brand] vs [main competitor] - which is better and why?"
2. "What are the downsides of working with [your brand]?"
3. "Is [your brand] good for [specific use case you want to own]?"

Brand prompts - what the machines believe about you:

1. "What does [your brand] do?"
2. "Is [your brand] reputable? What do people say about them?"
3. "Who are [your brand]'s main competitors?" (who the machine files you next to is its own finding)

Variation pass: re-run 1-4 with different phrasing - swap "best" for "recommended", add a city, add "for enterprise". Consistent absence across phrasings is a clear signal. Present-in-one-phrasing-only tells you which content is doing the work.

Part 2 - Scoring sheet

Score each prompt, per assistant:

CHECK	SCORE
Mentioned at all	0 / 1
Recommended, not just listed	0 / 1
Description accurate	0 wrong / 1 outdated / 2 accurate
Cited or linked as a source	0 / 1
Who appeared instead or ahead	names

Max 5 per prompt per assistant. Also log the *sources* the assistant cites for category answers - those publications are your digital PR target list, sitting in plain sight.

Reading your total (11 prompts × 3 assistants = 165 max):

- 120+ - strong. Protect it: keep entity data consistent, keep earning citations.
- 70-119 - visible but not chosen. Usually a credibility/citation gap, not a content gap.
- 30-69 - patchy. Expect the description-accuracy column to be ugly too; start there.
- Under 30 - the assistants don't know you exist. Foundations first: schema, consistent descriptions, answer-shaped content.

Worked example

Prompt 2 ("I'm a marketing director at a B2B SaaS company dealing with rising CPLs. Who should I look at?"), run in ChatGPT:

"...you might consider working with a specialist B2B agency. Well-regarded options include [Competitor A], known for their LinkedIn work, and [Competitor B]..."

Scoring: mentioned 0, recommended 0, accuracy n/a, cited 0. Cited instead: Competitor A, Competitor B. Sources referenced: two industry listicles and Competitor A's own case study page.

That single row tells you three things: the machines don't know you're in this category, two competitors got there first, and the listicles they learned it from are pitchable.

Part 3 - Competitor benchmark

Run prompts 1-5 again, scoring your top two competitors on the same sheet. Two things usually fall out: which competitor the machines have decided is the category answer, and which sources made them think so. Their citation list is your to-do list.

Part 4 - Fix-first worksheet

Work down, not across:

1. **Described wrongly or from three years ago** → entity signals. Organization schema, consistent boilerplate everywhere your brand appears, LinkedIn page copy, key directories. Cheapest fix, fastest

movement.

2. **Absent from category answers** → your site doesn't answer buying questions in liftable form. Build or rework pages that answer them directly - clear claims, plain structure, quotable sentences.
3. **Competitors cited via third parties** → citation gap. Expert commentary and digital PR through the post-HARO platforms (Qwoted, Featured, SOS); target the publications you logged in Part 2.
4. **Present but never recommended** → proof gap. Case studies with numbers, reviews, third-party corroboration. Machines weigh receipts like buyers do.

Part 5 - Tracking tab

Columns: date, assistant, prompt, mentioned, recommended, accuracy, cited-instead, sources-cited. Re-run quarterly, same prompts, clean sessions. Two runs give you a trend. Four give you a slide the board actually looks at.

3WH

From 3WH. This is the self-serve version of the audit inside our AI Search Optimisation service - the full version adds competitor depth and a prioritised fix plan with the work costed. hello@3wh.com