FOR SEO AND MARKETING AGENCIES · ONE-OFF WHITE-LABEL AUDIT

What you can put in front of a client

AI-Search Audit & Action Pack · two-page sales sample · $149 per client, paid once · you set the retail price

ANONYMISED SAMPLE. Real audit of a real US business, anonymised. The business did not commission or endorse it. “Client X” is a US heating, cooling, plumbing and electrical company in a major metro area; its name, location, contact details, competitors and identifying figures are redacted.

The agency’s commercial benefit. You receive a client-ready, white-label pack under your brand: a dated snapshot of what two official AI-search APIs answered, the sources checked, material inconsistencies on the client’s own website and listings, prioritised recommendations, and draft wording to fix them. You walk the client through it, decide what to implement, and charge what you choose. Our fee is your cost of goods — $149 for one full audit. There is no monthly minimum and no requirement to buy more than one client.

What this sample actually found

On one evening we asked ordinary homeowner questions — who to call for HVAC, AC repair, emergency plumbing, a new system, reviews, electrical work, a maintenance plan, and affordable AC replacement in [metro area] — through the official OpenAI and Perplexity search APIs.

The two engines did not agree. One almost never named this well-known local contractor; when it named a contractor, a local rival appeared more often, and the listing it cited most often was Angi. The other engine named Client X more often, but not consistently, even across rounds collected minutes apart. That is why this pack is a dated snapshot, not a ranking, and why a later change in those answers would not, by itself, prove that visibility had improved.

The useful work for an account team was on the company’s own site and listings. Every material finding below was checked against the live page on 23 November 2026 and still held on 24 November 2026.

What a client would see	Why it is useful to fix
HVAC rebate amounts do not agree across the website and the offers subdomain (and one line has a typo).	A homeowner — or an AI answer citing the site — can be shown three different figures for the same offer.
The homepage tile says “$[XX]/Month* New AC System” without “as low as”, credit-approval language, or the expiry that already appears on the offers subdomain ([expiry date]).	Easy to read as a cash price. The draft pack includes replacement wording and short FAQ lines.
Two Angi profiles for the same trading name: an approved page at 4.x from [n] reviews, and a “Not Angi Approved” page at 2.x from [m] reviews.	The engine that rarely named Client X cited Angi more than any other source.
Review badges, warranty lengths, phone numbers, an unsourced “#1” heading, and “open 24/7” versus appointment-dependent emergency language also disagree across pages and listings.	These are concrete edits an agency can take to a contractor, with sources attached.

IMPLEMENTATION-READY OUTPUT · EXTRACT FROM THE FULL PACK

A draft the agency can use, not just a finding

Every pack includes draft materials, clearly labelled as drafts for the agency and the client to approve. Nothing in this sample has been published. Here is the replacement homepage tile for the $[XX]/month offer (recommendation R2):

As currently worded on the homepage
"$[XX]/Month* New AC System" — low-APR financing is mentioned; "as low as", credit approval and the [expiry date] expiry are not.

Draft replacement
New AC system — payments as low as $[XX]/month for qualified applicants.
Subject to credit approval. Not a cash price for a complete system. May not be combined with other offers. Offer expires [expiry date]. Call [phone 1] or book online for terms.

The full pack also includes short FAQ lines to sit under that tile, a dated review-badge line, a factual replacement for the unsourced "#1" heading, and a LocalBusiness structured-data draft that uses only published offices and the primary website number. Recommendations are separated from work actually implemented: in this sample, nothing has been implemented.

What you get in the full audit

- A qualified two-engine snapshot from official OpenAI and Perplexity web-search APIs — not a signed-in consumer app, and not Google AI Overviews, Gemini or AI Mode.
- Verified sources: each material recommendation is tied to a page that resolves and actually supports the finding, with a quoted passage. Unsupported or inaccessible evidence is labelled, not dressed up as verified.
- Website and listing inconsistencies, prioritised recommendations, draft implementation materials, a client-ready summary, and a limitations statement.
- White-label presentation under your agency name (this sample uses a placeholder brand in the full pack).

What this is not. It does not measure Google rankings or consumer-app answers. It does not guarantee that fixing the issues will cause official APIs to name the client more often. It is not a monthly progress report. Mention rates can move between rounds on the same evening; a later movement would not by itself prove improvement or decline.

Method, in brief. Eight unbranded buyer questions, three rounds on each of two official APIs, collected on the evening of 23 November 2026 (48 of 48 answers returned). Website and listing checks were source-verified the same evening and re-checked on 24 November 2026 (23 of 23 material findings still verified, including a browser read of the BBB profile). Engines are reported separately and never blended into one score.

This sample is dated 24 November 2026. Offer-related findings were current when checked and are time-sensitive. The full 10-page white-label pack (also anonymised) is shown in a walkthrough, not sent as an unsolicited attachment.