

What your report looks like.

Sample — real estate agent, Southern California metro

14 prompts × 3 platforms · 42 live checks

THIS IS A SAMPLE, NOT A CLIENT REPORT

Every figure, finding and competitor below is **illustrative** — a composite of the patterns that come up again and again, written to show you the shape and depth of the deliverable. No real client's numbers, names, brokerage or market appear anywhere in this document, and none ever will. Your report runs on your own live data, and it belongs to you alone.

SECTION 01

The number

18%	OVERALL VISIBILITY	CHATGPT	GEMINI	CLAUDE
	15 points from a possible 84, across 42			
	live checks	21%	19%	14%

We ask the fourteen questions a real buyer or seller asks — in your city, in your specialty, at your price point — on three platforms, with no memory and no personalization. Every answer is scored: **2 if you are named**, **1 if your company is named without you**, **0 if you are absent**.

One number on its own means little. What matters is *which* questions produced it.

SECTION 02

The split that matters

ASKED ABOUT YOU BY NAME	SOMEBODY ACTUALLY SHOPPING
3 / 3	0 / 39

“Tell me about [you]. Are they any good?” All three platforms found this sample agent and described them accurately.

Best agent in the city. First-time buyer. Selling a home. Every neighborhood. Not named once.

Your reputation is usually not the problem. Being reachable by somebody who does not already know your name is the problem.

In this sample, six of those thirty-nine answers named the brokerage without naming the agent. The firm was in the conversation; the person was not in it with them.

SECTION 03

What we found, and why it decides the number

CRITICAL F1

The published record describes a business you no longer run

Asked to describe this agent, the platforms answered confidently and in detail — and named a set of neighborhoods with **no overlap** at all against the five areas the client actually wants. That single mismatch accounted for every zero on the neighborhood questions.

The machines were not ignoring them. They were repeating directory data years out of date, and recommending them into the wrong part of the county.

FIX Correct the service areas on every profile the client controls. Free, and usually an afternoon.

CRITICAL F2

No machine-readable facts on the website

Zero structured data — none of the hidden markup that states a name, licence, service area and specialty as *facts* rather than prose. With nothing of their own to read, the only machine-readable claims about this agent were the stale third-party ones in F1.

The page description compounded it: warm, personal, and containing no city, no neighborhood, no specialty and no credential. It would have fitted any agent in the country.

FIX Agent schema naming the real service areas and credentials, plus a rewritten title and description.

HIGH F3

Credentials held, and invisible

This sample agent held four industry designations. **None appeared anywhere a machine could read.** In the same run, one platform told the buyer to prefer agents holding a specific certification — one the client already had, and published nowhere.

FIX State every credential with its issuing body, in page copy and in structured data.

HIGH F4

Reviews concentrated on one platform

Strong reviews, almost all on a single site, with a low single-digit count on Google — the source the AI answers lean on hardest. Corroboration across platforms is what these systems weight; the same reviews spread across four sites read as stronger than all of them in one place.

FIX A structured review campaign from the client's own past-client list.

Published under several different names

Three name variants across their own site, their licence and a major directory, with nothing connecting them. One platform reconciled it mid-sentence; a system with less patience treats them as different people and splits their credibility.

FIX Settle on one public name, correct the outliers, and bind every profile together in structured data.

SECTION 04

Who gets recommended instead

Across the questions where this agent was not named, we count who was — by number of distinct answers, not total mentions.

Illustrative — your report names the real firms and people in your market

WHO	ANSWERS NAMING THEM
Competitor Team A	13
Competitor Team B	9
Competitor Agent C	7
Competitor Firm D	6
Competitor Agent E · Competitor Firm F	4 each
The competitor the client expected to see	1

That last row comes up often. Clients usually name a rival they watch closely, and the audit finds that rival is not who is actually winning the answers. Knowing the real list changes what you build.

SECTION 05

What fixing it involves

Every report splits the work in two, because the split matters more than the list.

YOURS · FREE

What you can do yourself

Correct your service areas everywhere you can edit. Ask past clients for Google reviews. Use one consistent name. Claim the directory profiles you are missing from.

We hand this over in full, in writing, whether or not you hire us for anything else. It is also the work most likely to move the number.

WE HANDLE

The plumbing

Structured data. Rewritten titles and descriptions. Google Search Console, Bing and instant indexing. Location and specialty pages. Review-request templates so your campaign is copy-and-paste.

Worth being straight about the indexing work: it makes you *findable* faster. It does not make you *recommended*. Those are different problems and the first is the easier one.

SECTION 06

How you know whether it worked

The number in Section 01 is a baseline, and that is its whole purpose. After the fixes, we re-run the **identical** questions on the same platforms with the same scoring, and compare. Same test, before and after.

Indexing changes show up in weeks. Being *recommended* takes longer, and we will not put a date on it. These systems are not deterministic — we improve the inputs they read and track the trend. Nobody can guarantee a specific answer, and anyone who tells you otherwise is guessing.

You keep the baseline either way, whatever you decide to do next.

About this document. A sample of the AI Visibility Audit deliverable. All figures, findings and competitor labels are illustrative composites shown to demonstrate format and depth. No client's data, name, market or brokerage appears here. Client reports are confidential and are never shown to anyone else.

Method in a real report. 14–16 prompts × 3 platforms via API with web search enabled, no memory or personalization. Scoring: 2 = named, 1 = company named without you, 0 = absent. Entity findings verified against live pages on the day of the run; where a page refuses our request, the report says so and draws no conclusion from it.

Haas Business Solutions • Dalan Haas • dalan@haasbusinesssolutions.com • 808-483-0960