

The report you hand the person who asked.

This is a complete sample of the client-ready AI-visibility report — every section the client or boss will actually read, filled with clearly-invented example data so you can judge the shape before you measure anything. Each section is annotated with the question it answers, because that's why it earned its page.

Built for the most common job in this discipline: *"my client wants to know what ChatGPT says about them — how do I report that?"* Inside the product, this document is generated from your runs and shared as a no-login link.

[Skip to content](#)

AI VISIBILITY REPORT · ILLUSTRATIVE SAMPLE

Northwind CRM — August 2026

Runs #1–6 · Mar 12 – Aug 4, 2026 · 48 buyer prompts × 4 engines, re-asked fortnightly · 384 answers per run

SAMPLE — EVERY NUMBER IS INVENTED

1 · WHERE YOU STAND

The four numbers, with their error bars

The question it answers: "So — are we visible or not?" Four rates, each with its 95% confidence interval, so a client can tell a real change from noise without trusting you on it.

MENTION RATE

41% +14 ✓

95% CI 36–46% · beyond noise vs run #1

SHARE OF VOICE

24% +9 ✓

95% CI 20–28% · #2 of 6 tracked

RECOMMENDED

18% +4 ~

95% CI 14–22% · within noise

AVG RANK

2.3


when listed · lower is better

✓ marks movement whose confidence intervals no longer overlap the baseline run. The +4 on Recommended stays marked ~ — it might be real, but the data can't prove it yet, and the report says so rather than rounding up.

2 · WHAT THEY SAY

How engines frame Northwind when they name it

The question it answers: "What does ChatGPT actually say about us?" — the question that started the engagement. The split, plus the recurring narratives with real example lines.



"For a small agency that lives in its inbox, Northwind is usually the easiest of the three to actually roll out."

Recurring narrative — 11 of 24 mentioning answers · illustrative


"Northwind covers the basics but has fewer native integrations than the bigger suites."

The negative 9% — one consistent objection, which is fixable content, not fate · illustrative

3 · AGAINST WHOM

Who the engines name in your category

The question it answers: "Who's getting the deals we don't see?" How often each tracked brand appears across the same 384 answers — the shortlist your buyer is actually shown.



Run #1 had Northwind fourth. The climb to #2 is the headline a client remembers — and it's defensible because the same 48 prompts were asked every run.

4 · ENGINE BY ENGINE

Where you're strong, and where you're absent

The question it answers: "Which assistant should worry us?" The same brand can dominate Perplexity and barely exist on Gemini — the average hides the gap the next quarter's work should target.

[Skip to content](#)

ENGINE	MENTION RATE	95% CI	TREND OVER 6 RUNS
Perplexity	52%	45–59%	rising — strongest surface
ChatGPT	44%	37–51%	rising since the pricing rewrite
Gemini	33%	27–39%	flat — the gap to work
Claude	28%	22–34%	flat, positive when present

5 • THE EVIDENCE TRAIL

Where the engines get their facts

The question it answers: "Fine — what do we *do* about it?" The domains the engines cited when answering, and whether each one mentions you. A trusted source that names only your rivals is the to-do list.

CITED DOMAIN	SHARE OF ANSWERS CITING IT	NAMES NORTHWIND?
g2.com	31%	YES
reddit.com	24%	NO — NAMES BELLROSE & KEYSTONE
capterra.com	17%	YES
wikipedia.org	12%	CATEGORY PAGE ONLY

The reddit row is where the next retainer month goes: the engines already trust it, and it currently tells them about two rivals.

6 • ACCURACY

What the engines get wrong about you

The question it answers: "Is anything they say actually damaging?" Being described wrongly can cost more than being absent — a wrong price or a

"no API" claim quietly disqualifies you from shortlists.

CLAIM FOUND IN ANSWERS	REALITY	STATUS
"Northwind has no public API"	REST API shipped Jan 2026	RESOLVED — DOCS PAGE PUBLISHED RUN #4, CLAIM GONE BY RUN #6
"Starts at \$49 per seat"	\$29 per seat since March	OPEN — PRICING PAGE UPDATED, ENGINES STILL CITING AN OLD REVIEW

7 · PROOF OF WORK

What we shipped, and what it measurably moved

The question it answers: "What am I paying you for?" Each shipped change is recorded against a baseline run and a treatment run, and the lift is only claimed when it clears the confidence intervals. This is the section that renews retainers.

CHANGE SHIPPED	RUNS COMPARED	MEASURED EFFECT
Rewrote /pricing answer-first, added Product JSON-LD	#2 → #4	Mention rate +9 pts ✓ beyond noise
Published the Northwind vs Bellrose comparison page	#4 → #6	Share of voice +4 pts ~ within noise — holding, re-measuring

8 · NEXT

The three moves for the next period

The question it answers: "What happens next month?" Ranked by expected effect on the numbers above, each tied to the finding that justifies it.

Skip to content

1. **Get described on the reddit threads the engines cite** — the one trusted source that names only rivals (section 5).
 2. **Close the Gemini gap** — publish the integration answer pages Gemini's answers keep reaching for (section 4).
 3. **Kill the stale \$49 price** — request re-crawl of the outdated review, keep the corrected pricing page canonical (section 6).
-

APPENDIX

Method — why these numbers hold up

The question it answers: "How do I know this isn't made up?" The four questions to ask any vendor: how many runs, what dates, what prompt set, what spread. A report that can't answer them is a screenshot.

- The same 48 buyer prompts asked on every run — the suite is snapshotted, so a moved metric is never a moved questionnaire.
- 6 runs, fortnightly, Mar 12 – Aug 4 2026 · 384 answers per run across 4 engines.
- Every rate carries a 95% confidence interval; movement is claimed only when intervals don't overlap.
- Full method, limitations included, published at mentionbeat.com/kb/methodology (mentionbeat.com/kb/methodology.html).

What this sample cannot show you

Your numbers. Everything above is invented to show the shape — the real report is generated from measured runs on your own prompts, and every figure arrives with the confidence interval those runs earned. The free checker won't produce this document, but it will answer the first question in it — whether you're in the answer at all — in about a minute: run it on **your site** (mentionbeat.com/checker).

Inside the product (mentionbeat.com/product), this report opens from the dashboard in one click, shares as a no-login link for the person who asked, and can arrive by email on a schedule.