

Client AI Search Audit

Placeholder-safe client report template for AI Overviews, Gemini, ChatGPT, AI Mode, and Perplexity visibility.

Prepared for: Example client. Scope: example.com priority pages. Evidence collected 2026-05-22; replace placeholders only where source IDs corroborate the claim.

Audit manifest

Prepared for: Example client

Scope: example.com

Period: 2026-04-01 → 2026-05-22

Audited by: Marcus Quinn

Sources: 27 captured

Next audit: 2026-W30

Executive summary

Example client appears in some answer-engine responses for category and comparison prompts, but source coverage is uneven. Priority gaps are retrieval eligibility, source-card proximity, and third-party corroboration parity. **EVIDENCE:** PARTIAL

3/5

Engines with at least partial visibility. Source: C001. Trend: +1.

12

Priority URLs reviewed against retrieval, source-card, and parity checks. Source: C002, C003.

27

Distinct evidence references captured this window. Source: ledger. Trend: +9.

6

P0/P1/P2 roadmap items sized for the next monthly cycle. Source: §04.

DECISION: prioritise comparison and pricing pages before lower-impact article refreshes. Owner: Editorial · 2026-W23.

▲ Action Prompt

COPYABLE ACTION PROMPT

Reference this report action: Decision: prioritise comparison and pricing pages before lower-impact article refreshes. Owner: Editorial · 2026-W23.

Guide me through the tools, resources, accounts, permissions, source material, and access needed to take this action. Break the work into numbered steps, call out any missing inputs before execution, include safe handling for credentials or confidential data, and finish with verification evidence I can capture.

§ 02 | Engine
findings | 5 engines ·
5 rows

§ 03 | Source ledger
| 27 captured ·
4 shown

§ 04 | Priorities &
weighted findings |
2 P0 · 1 P1 · 2 P2 ·
1 done

§ 05 | Preserve & fix
| 3 keep · 4 fix

§ 06 |
Implementation
brief | Engineering
handoff

§ 07 | Verification
checklist | 5 items ·
2 done

CONTENTS

Client AI Search Audit

Audit manifest

Executive summary

1. Engine findings

2. Source ledger

3. Priorities & weighted findings

3.1 P0: Comparison page retrieval gap

3.2 P0: Perplexity citation gap on priority prompts

3.3 P1: Pricing facts lack nearby evidence

3.4 P2: ChatGPT routes recommendation via third-party profiles

3.5 P2: AI Mode direct-answer opening is too soft

3.6 P3: Long-tail category articles need a refresh

3.7 Done: Source-card component shipped to comparison template

4. Preserve & fix

4.1 Preserve

4.2 Fix

5. Implementation brief

5.1 Engineering handoff

5.2 Editorial note


6. Verification checklist

CHAPTER 1

1. Engine findings

Five surfaces, one priority prompt set. Coverage is bimodal: strong on Google-family surfaces, weak elsewhere. Perplexity is the only engine with zero citation in this window.

| Engine | Finding | Evidence | Next action |
|-------------------------|--|---------------------------|---|
| AI Overviews – Google | Brand cited for two comparison prompts. | EVIDENCE: VERIFIED | Add source cards to pages with citations but weak snippets. |
| Gemini – Google | Mentions brand but misses pricing constraints. | EVIDENCE: PARTIAL | Add visible plan constraints and updated date. |
| ChatGPT – OpenAI | Recommendation inferred from third-party profiles. | EVIDENCE: INFERRED | Improve profile parity and owned corroboration. |
| AI Mode – Google | Partial feature coverage, no comparison citation. | EVIDENCE: PARTIAL | Strengthen direct-answer opening on comparison page. |
| Perplexity – Perplexity | No citation found in priority prompt set. | EVIDENCE: MISSING | Build third-party corroboration and rerun prompt set. |



Coverage is bimodal. Google-family surfaces are above 50% on most prompts; the rest sit below 45%. Perplexity is the outlier and should be re-tested after corroboration work ships.

CHAPTER 2

2. Source ledger

Every material claim links back to a source ID. The full ledger of 27 entries lives in the secure evidence folder; only redacted IDs are reproduced here.

C001 | Prompt
capture batch |
High
confidence ·
raw transcripts
plus
screenshots in
secure
evidence
folder.

C002 |
Rendered
crawl | High
confidence ·
confirms
client-
rendered
sections
missing from
first fetch.

C003 |
Analytics
export |
Medium
confidence ·
prioritises
pages with
commercial
intent.

C004 | Third-
party parity
review |
Medium
confidence ·
category and

pricing facts
mismatch
owned page.

**PRIVATE EVIDENCE
RULE**

Public issues, PRs, and examples must not include non-public names, restricted URLs, local paths, screenshots, or raw exports. Use placeholders and keep evidence in the approved secure storage location.

3. Priorities & weighted findings

Each item is tied to state colour and source IDs: critical regressions in red, watch items in amber, informational items in blue, completed work in green.

3.1 P0: Comparison page retrieval gap

Comparison content answers the right intent but key criteria are below client-rendered sections. Move the answer, feature table, pricing caveats, and source IDs into crawlable first-fetch

HTML. EVIDENCE: **VERIFIED**

Owner: Editorial. Due: 2026-W23. Sources: C001, C002.

3.2 P0: Perplexity citation gap on priority prompts

No Perplexity citations in the priority prompt set. Build a corroboration pass across the three highest-authority third-party profiles, then rerun the prompt set under fresh conditions. EVIDENCE: **MISSING**

Owner: PR. Due: 2026-W24. Sources: C001, C004.

3.3 P1: Pricing facts lack nearby evidence

The pricing page states plan limits but lacks updated date, source card, and consistent third-party parity. Add canonical values and cite source IDs near the claims. EVIDENCE: **PARTIAL**

Owner: Pricing PM. Due: 2026-W25. Sources: C003, C004.

3.4 P2: ChatGPT routes recommendation via third-party profiles

Owned pages are rarely the source for ChatGPT recommendations in this category. Improve profile parity across the three highest-traffic review sites and add owned corroboration anchors. EVIDENCE: **INFERRED**

Owner: PR. Due: 2026-W28. Sources: C004, C019.

3.5 P2: AI Mode direct-answer opening is too soft

The comparison page opens with hedge copy that AI Mode does not lift as a direct answer. Rewrite the opening sentence as a single-claim, source-anchored statement. EVIDENCE: **PARTIAL**

Owner: Editorial. Due: 2026-W29. Sources: C001, C012.

3.6 P3: Long-tail category articles need a refresh

Older category articles still rank but draw little attention from answer engines. Schedule a refresh after the comparison and pricing work has shipped and been verified.

Owner: Editorial. Due: 2026-Q4. Source: C011.

done

3.7 Done: Source-card component shipped to comparison template

The reusable source-card component is now embedded in the comparison page template. First-fetch HTML now carries citation anchors and updated dates next to every claim.

Owner: Engineering. Verified: C001. Next: re-run prompt set.

4. Preserve & fix

4.1 Preserve

- Clear category positioning across owned and earned surfaces.
- Strong customer-proof signals on the comparison page.
- Comparison-intent coverage is relevant and consistent.

4.2 Fix

- Critical content hidden below client-rendered sections.
- Unsupported superlatives in feature copy.
- Stale third-party profile facts across three profiles.
- Perplexity source-coverage gap.

5. Implementation brief

Task: Improve `/compare/example-vs-competitor` for AI search citation readiness.

Files: comparison page template, pricing facts component, source-card component.

Acceptance: raw HTML contains direct answer, source IDs C001–C004, updated date, and criteria table.

Verification: rerun AIO, Gemini, ChatGPT, AI Mode, and Perplexity prompt set separately.

Owner: Editorial + Engineering pair. **Due:** 2026–W23. **Rollback:** feature flag `cmp-v2` reverts to current template.

5.2 Editorial note

Keep the criteria table inside the first viewport on desktop and inside the first scroll on mobile. The current placement is the single largest contributor to retrieval failure on this surface.

6. Verification checklist

- Confirm live evidence collection date and source IDs.
- Re-run crawl after implementation and compare raw vs rendered HTML.
- Re-run per-engine prompt set and record each answer separately.
- Update third-party profile parity table.
- Convert recurring monitoring into a monthly report routine.

NEXT: re-audit fires 2026-W30. Calendar invite attached to the secure handoff folder. Owner: Ops · weekly cadence.

▲ Action Prompt

COPYABLE ACTION PROMPT



Reference this report action: Next: re-audit fires 2026-W30. Calendar invite attached to the secure handoff folder. Owner: Ops · weekly cadence.

Guide me through the tools, resources, accounts, permissions, source material, and access needed to take this action. Break the work into numbered steps, call out any missing inputs before execution, include safe handling for credentials or confidential data, and finish with verification evidence I can capture.

