

Sample SEO & GEO Audit — Anonymised

Prepared by Tasdid Hossain Siddique · Subject: a US business-to-business services client (third-party logistics), ~30 indexed content pages

Anonymised at the client's level: no client name, URLs or commercial data. Findings, figures and the method are as produced.

HOW I AUDIT

I run every page through a four-stage checklist and log the result per URL, so nothing is diagnosed twice and no page is re-audited from scratch:

STAGE	WHAT IT COVERS
0.x Pre-flight	Can the page actually be reached and read — crawler access by user agent, robots and sitemap state, indexation, redirect integrity. Nothing downstream is trustworthy until these pass.
1.x SEO	Intent match, title and meta, heading hierarchy, internal linking, Core Web Vitals, cannibalisation against sibling pages.
2.x AEO	Answer-first structure, question-shaped headings, inline source attribution, FAQ and entity markup as <i>served</i> , extractability of each section as a standalone answer.
3.x GEO	Presence and absence in AI Overviews and assistant answers for the target query set, citation and mention counts over time, entity consistency across the web, competitor citation share.

The rule that governs the whole document: a finding is only recorded against the **rendered, served page**, never against the template or a validator's verdict. Three of the findings below exist only because of that rule.

HEADLINE FINDINGS

P1 FAQ structured data validated on 29 pages and shipped empty on 26 of them

FAQPage markup passed Google's Rich Results Test across the content library. On the served HTML it was present on three pages. The markup lived in the template layer and never reached the rendered output. Every report that counted "FAQ schema: present" was counting the template.

Why it matters for GEO: question–answer pairs are among the most directly extractable structures for an answer engine. The client had been told this was done for over a year.

Fix: diff template output against served HTML as a standing check; rebuild the markup at render. **Effort:** low. **Confidence:** verified.

P1 27 backlinks across 13 referring domains pointing at a 404

A retired URL had kept acquiring links. No redirect existed. Recovered with a single redirect rule and **no new content** — the highest return-per-hour item in the engagement.

P2 Duplicate meta description emitted by the theme

The theme printed a second description tag whenever a post carried a manual excerpt. Diagnosed to the theme rather than the SEO plugin, so the fix was to clear the excerpt field rather than modify the theme — a reversible change with no upgrade debt.

P2 Crawler access confirmed clean — and this needed checking before anything else

Googlebot, Bingbot, GPTBot, ClaudeBot, PerplexityBot, CCBot, OAI-SearchBot and Google-Extended all returned 200 at the user-agent layer. The CDN was not blocking AI crawlers, which is the first thing to rule out on a GEO engagement and is frequently assumed rather than tested. Noted as a limitation: verified at the user-agent layer only, not by crawler IP or ASN.

A generated 11ms.txt was found to be plugin boilerplate with no curatorial value — flagged as cosmetic, not a ranking factor, and deprioritised rather than "fixed".

THE JUDGEMENT CALL THE AUDIT ACTUALLY TURNED ON

All four audited pillar and cluster terms returned an AI Overview with the client absent from it — and absent from the organic top ten — despite on-page work that graded well on every SEO and AEO check.

The honest reading is that **the binding constraint was site authority, not content quality**. Page authority sat at 18 with effectively no external links into the audited pages, against roughly 166 unresolved 404s and inconsistent name/address data across third-party business directories.

So the recommendation was to stop producing content passes. Priority order became: resolve the 404 inventory and reclaim the links already pointing at them, correct the directory and citation inconsistencies, then build authority — and only then return to content. Recommending less of the work the client was already paying for is uncomfortable, but another well-optimised page would not have moved the AI Overview.

WHAT PROVED THE METHOD, ON A SINGLE PAGE

One commercial page was rebuilt against two rules only:

- **Every heading answers its own question in the first sentence** — no wind-up before the answer.
- **Every statistic names its source inside the same sentence**, not in a reference block at the foot of the page.

Result: the page moved from roughly position 51 to 20, and **earned a Google AI Overview citation** for its core query. The ranking movement was secondary; the citation was the signal, because it meant an answer engine had begun treating the page as a source rather than a result. That page became the template for everything published afterwards.

MEASUREMENT, AND WHERE IT IS STILL WEAK

SIGNAL	OBSERVATION
Organic sessions	~99 → 294 per month (+197%) across 455 ranking keywords, following a pillar-and-cluster rebuild mapped to search intent
AI citations	204, down from a peak near 280 — the decline correlates with a publishing gap (4 pieces in one month with a 20-day hole, against 9 in each of the two preceding months) rather than with any algorithm event
Mentions vs citations	1 mention against 204 citations — the brand is being <i>used</i> as a source far more often than it is being <i>named</i> . Treated as the central open problem, not a vanity gap
Share of impressions in AI surfaces	~20% of impressions now appear in Google's AI features report while average position is largely unchanged — the metric the client reports on is drifting away from where discovery happens

Stated limitation. This tracking is manual plus a general SEO suite's AI-visibility view. It is enough to see direction and to correlate a decline with a publishing gap. It is **not** enough to attribute movement to specific prompts or specific engines, and I said so in the report rather than implying more precision than the data carried. A dedicated AI-visibility platform is the correct instrument and I would want one.

HOW THE WORK WAS PRODUCED

The per-URL audit is run semi-automatically: an agent pulls Search Console data, crawls the page, compares template output against served HTML, and appends findings to a running log keyed by URL. Two early failures shaped the controls now in place — it reported structured data as present because the template contained it, and it produced plausible statistics that did not exist. The resulting rules: diff template against render, and

every figure must resolve to a named, dated source before it enters a report. Volume is automated; whether a claim is true is not.

Tasdid Hossain Siddique · Chattogram, Bangladesh (UTC+6) · tasdidsiddique@gmail.com · [linkedin.com/in/tasdid-hossain-siddiqui-172863413](https://www.linkedin.com/in/tasdid-hossain-siddiqui-172863413)