

BEFORE I RUN THIS ON YOUR SITE, HERE'S WHAT IT FOUND ON MINE.

Site: marikaolson.com · my own site · audited June 2 and June 15, 2026 · every finding verified against the then-live site on its original host

Marika Olson Consulting · operations & website design for small business owners · Pacific Northwest

AI VISIBILITY SCORE: 54 / 100

FAIR · NEEDS WORK, MOVING UP



This is my own site, scored on the exact framework I would use on yours, captured the week I moved it off the platform I used to use. I am not going to show you a polished client win and call it a sample. I will show you mine, gaps and all, because the honest version is the point. **The headline: the site is well built and well written, and it still barely registers with AI search.** That is not a plumbing failure. It is a content and authority gap, and it is the most common pattern I see on small business sites.

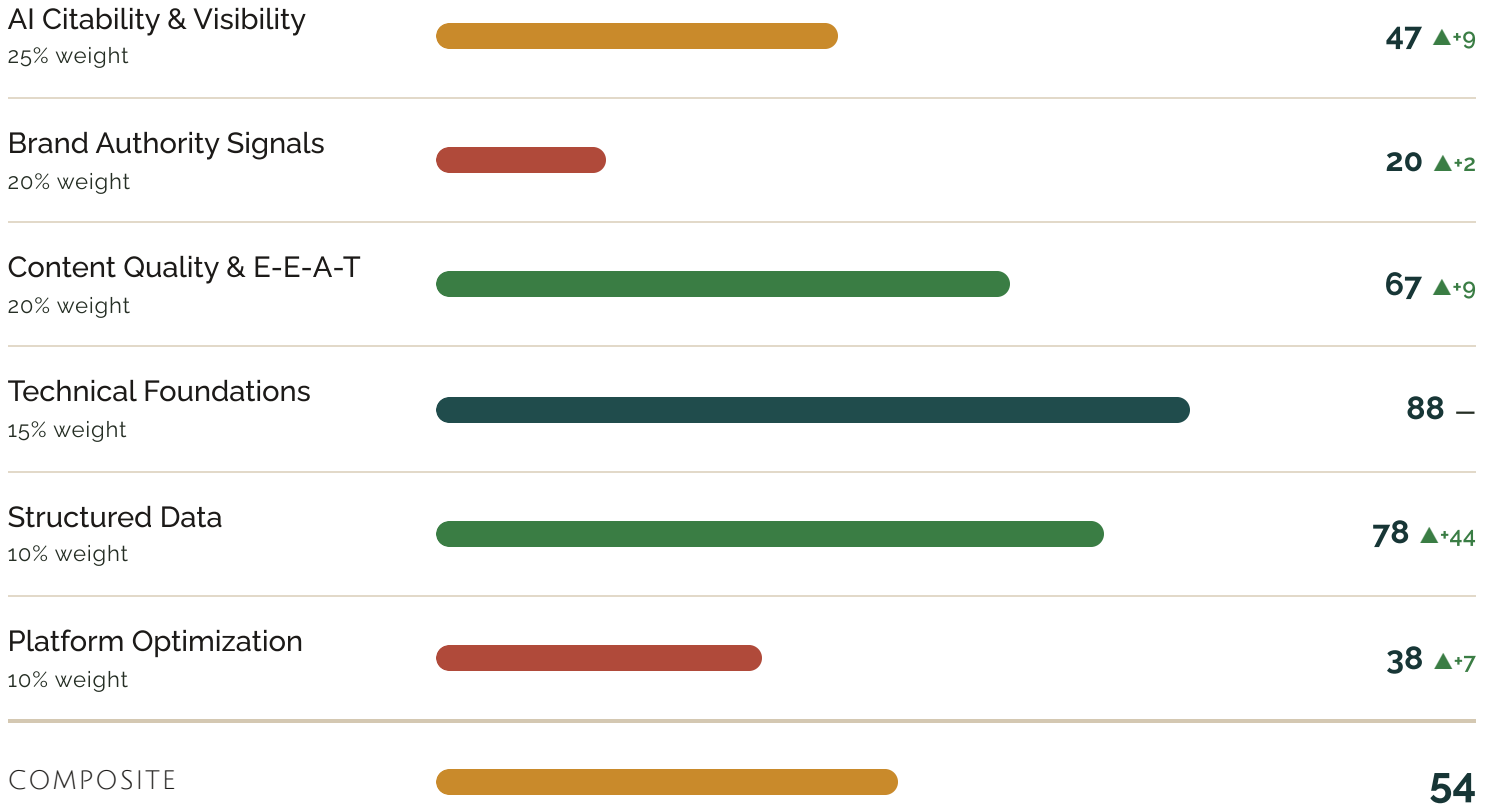
THE REASON I TRUST THE NUMBER



I ran this same audit on myself on **June 2** and scored a **46**. Then I started fixing the quick wins. Thirteen days later it is a **54**. The biggest single jump was structured data, which went from 34 to 78 once I repaired a broken schema block. The method tracks real movement. That is the whole reason to measure instead of guess.

WHERE THE SCORE COMES FROM

SIX WEIGHTED DIMENSIONS



Scale: 0–20 Critical · 21–40 Poor · **41–60 Fair** · 61–80 Good · 81–100 Excellent. Green ▲ = change since the June 2 baseline.

AI CITABILITY & VISIBILITY 47 / 100

The content that exists is genuinely citable. My About bio and a new post on AI visibility are clean, self-contained answer blocks, the kind an engine quotes directly. What caps the score is volume, not quality.

Working in my favor:

- **AI crawlers are fully allowed.** GPTBot, ClaudeBot, PerplexityBot, Google-Extended and the rest can all reach every page. Access is not the problem.
- **The few strong pages punch above their weight** because the writing is dense and factual.

The drag:

- **Only three blog posts.** There is almost no answer surface for the questions a buyer actually asks an AI. You cannot be cited on topics you have not written about.
- **No 11ms.txt**. The platform I used to use couldn't host one at the root, so the content map AI engines look for simply wasn't there. It was one of the factors that tipped me toward moving to a build I fully control.

BRAND AUTHORITY SIGNALS 20 / 100

My weakest dimension, and the most honest thing in this report. On a real buyer query like "fractional operations consultant on the Olympic Peninsula," my own site does not surface. Larger firms own that result. Here is exactly why.

- **No Wikipedia or Wikidata record.** That is the strongest entity signal AI trusts, and it is a hard zero. A solo consultancy will not merit a Wikipedia article, so the realistic path is a Wikidata entry plus regional mentions.
- **No third-party footprint.** Nothing on Reddit, YouTube, or review platforms. Every signal pointing at me is one I created myself.
- **The entities I do have are now wired but thin.** LinkedIn, Instagram, and a claimed Google Business Profile are connected through schema, but the Business Profile only had one review on it at the time of the audit. (*Perhaps you'd be so kind as to add a positive review after reading this audit?*) The wiring is right; the destinations need filling.

CONTENT QUALITY & E-E-A-T 67 / 100

E-E-A-T stands for Experience, Expertise, Authoritativeness, and Trust, the four things both Google and AI engines weigh when they decide whether to rely on a source. My experience and expertise score high. Real case studies with dollar figures, a 17-year operations spine, named credentials.

- **The service pages were rebuilt deep and FAQ-backed** recently, which is most of the recent gain. Those FAQs are the single most AI-citable asset on the site.
- **The ceiling is authority and volume.** One testimonial and three posts. Adding a privacy policy and terms page would also pick up cheap trust points I am currently leaving on the table.

TECHNICAL FOUNDATIONS 88 / 100

The healthiest dimension, and the one that surprises people. The site is fast enough, mobile-clean, and the content is server-rendered, which means AI crawlers that do not run JavaScript still see the full text on every page. Nothing important is hidden behind code.

- **The real gaps were platform-imposed.** Three security headers and an editable robots file were things the platform I used to use would not let me add. A static build fixes all of them.
- **One small JavaScript blind spot:** the booking widget on the contact page is invisible to non-JavaScript crawlers. Low impact, easy fix on the next build.

STRUCTURED DATA 78 / 100

The big turnaround, from 34 to 78. Structured data is the machine-readable summary of who you are that lives invisibly in the page code and tells AI how to understand your business. Mine was broken two weeks ago; a single syntax error was invalidating the whole block.

- **It is now valid and server-rendered:** a connected Organization, Person, and Service graph, plus FAQ markup on the service pages, all readable by crawlers that skip JavaScript.
- **The remaining ceiling was platform-gated** (speakable markup, per-page schema control). That headroom, into the low 90s, is exactly what a build I control unlocks.

PLATFORM OPTIMIZATION 38 / 100

How ready the site is to be surfaced by each specific engine: Google AI Overviews, ChatGPT, Perplexity, Gemini, and Bing Copilot. Crawlers are welcome everywhere, so access is not the limiter. Two things are.

- **Content volume**, again. There is not enough surface for any engine to treat the site as a primary source on a topic. That does not mean writing AI slop. It means crafting evergreen content that answers the questions my potential clients are actually asking. This is a marathon, not a sprint.
- **Entity confusion**. My name still bleeds into an earlier public-sector identity in the sources these engines read, so the consulting business has no clean record to attach to. A Wikidata entry would disambiguate it.

THE PLAN I AM ACTUALLY RUNNING ON MYSELF

- 1** **HIGH** **Publish on a real cadence and build query-shaped answer pages.** Content volume is the number one lever, and nothing else compounds without it. Write the pages that answer what buyers actually ask: what a fractional ops partner costs, when to hire one, how to make a business run without you.
- 2** **HIGH** **Build the off-site authority I am missing.** Create a Wikidata entry to anchor and disambiguate the business, grow real reviews on the Google Business Profile, and turn my best testimonial into a full before-and-after case study page.
- 3** **HIGH** **Move to a build I fully control.** Per-page schema, security headers, an owned robots file, instant re-indexing, a clean root-level `11ms.txt` : none of it was tunable on the platform I used to use. So I moved to a static build for exactly this reason. (Yes, I took my own advice: this report is hosted on it.)
- 4** **MED** **Add FAQ blocks with FAQ schema across the site,** plus visible dates. This is the fastest way to feed AI answer boxes directly.
- 5** **MED** **Add a privacy policy and terms page.** Cheap, quick trust points that the audit is currently docking me for.
- 6** **MED** **Ship security headers and instant re-indexing on the new build,** both of which the current platform blocks.

Where not to spend a minute

The AI-crawler block toggle people worry about was already set to allow on my site, so there was nothing to fix there. Security-header and heading-semantics tweaks on the platform I used to use were cosmetic, near-zero impact. I would rather you spend the time on content.

WHAT A FULL ENGAGEMENT ADDS

A ONE-TIME SCORE IS A SNAPSHOT. THE PROGRAM PROVES THE FIXES WORKED.

1 · AI-visibility polling, the rigor layer. We define 15 to 25 buyer-intent prompts, run them across ChatGPT, Perplexity, Gemini, and Google AI Overviews on a monthly cadence, and report citation rate, trend over time, and share of voice against your named competitors. AI answers are volatile, so a single check is noise. Repeated waves are the only honest measure.

2 · Before-and-after attribution, which I am living right now. Baseline 46, then 54 after the first fixes, then the move to a build I control. The re-audit on the new build is the proof the work moved the number. You get the same on your site: a baseline, the fixes, and a 30, 60, and 90-day re-audit that shows the deltas. And the whole thing is anchored to the one or two business KPIs that matter most to you, agreed up front and chosen so the work can show real cause and effect, not just a score that moved.

Plus full-catalog per-page citability scoring, a complete schema matrix with ready-to-paste code, competitor benchmarking, an E-E-A-T scorecard, and the technical fixes a static build unlocks.

IF I WILL SHOW YOU MY OWN NUMBER, GAPS AND ALL, YOU CAN TRUST THE ONE I GIVE YOU ON YOURS.

Method. Scored on MOC's standard GEO rubric across six weighted dimensions (AI citability 25%, brand authority 20%, content and E-E-A-T 20%, technical 15%, structured data 10%, platform readiness 10%), synthesized from five specialist passes and verified against the live site on June 15, 2026. The June 2 baseline of 46 was run the same way, which is what makes the movement comparable. Brand-mention checks use the public web plus the Wikipedia and Wikidata APIs, where absence is the metric that matters, because not being findable is the finding.

Prepared by Marika Olson Consulting · marikaolson.com · This is a real audit of my own site, shared as a sample on purpose.