

The GEO Checklist Deep-Dive: 12 AI Visibility Fixes, Fully Operational

Most AI-visibility wins are structural, not creative. This is the full operational version of our 12-fix checklist: what each fix is, exactly how to ship it, and the citation scoreboard that tells you whether any of it is working.

We published a 12-fix GEO checklist built on one finding that repeats across every program we run: most of the AI-visibility opportunity is structural, and it is being ignored. The checklist told you what to do. This deep-dive is the operational version: for each fix, exactly how to ship it, what to check, and how to know it worked.

The order matters. The list is ranked by **effort-to-impact ratio**, not importance. The first items take hours. The last items take weeks. Start at the top. Don't skip to #12.

One honest caveat before anything else: none of these fixes will move citation share if you have nothing worth being cited for. Fix the structure, yes, but also have a point of view. Structure makes credibility findable. It does not manufacture it.

WHO THIS IS FOR

Who this is for

Founders, operators, and content leads who want their business named and cited when someone asks ChatGPT, Perplexity, Gemini, or Google's AI a question in their space. You need edit access to your website and a couple of focused blocks of time. You do not need an engineering team for the first eight fixes.

Phase 1: The quick wins (ship today)

Three fixes, a few hours total, and they remove the blockers that make everything else pointless.

Fix 1: Add llms.txt at your root

Create a plain-text file served at `/llms.txt` (for most stacks, drop it in your public folder). It should cover three things:

- **Who you are.** One or two sentences, the same description you use everywhere else.
- **What you are authoritative on.** The two or three topics you genuinely know.
- **Where the canonical content lives.** A short list of your most important URLs with a line on what each covers.

This is the emerging standard for telling AI agents what your site is, in a format built for them rather than scraped from your nav. It costs about an hour, and it gives engines a canonical, hallucination-resistant summary of your business instead of leaving them to guess. Early, low effort, compounds.

How to verify: open `yourdomain.com/llms.txt` in a browser. If it loads as plain text and a stranger could understand your business from it, done.

While you are at your root, audit `robots.txt` too. If GPTBot, ClaudeBot, PerplexityBot, or Google-Extended is blocked there, nothing else in this guide matters: engines cannot quote what they cannot read. And check one level deeper, because a CDN or firewall rule can silently block those crawlers even when robots.txt allows them.

Fix 2: Put a visible "Last Updated" date on every key page

Generative engines weight freshness heavily for any query with a temporal component, which is most commercial queries. A page dated 2024 loses to the same page dated 2026, all else equal.

Two rules make this real instead of cosmetic. First, the date must be visible on the page, not only in metadata: engines and humans both read it. Second, only update the date when you actually touch the content. A "last updated" stamp that auto-bumps on every deploy is the kind of synthetic signal that erodes trust in everything else on the page.

How to verify: open your five most important pages as a visitor. Can you find a date within a few seconds? Is it accurate?

Fix 3: Add a named author to every page

A real name, linked to a real bio, with a consistent identity across the site. "Staff" and "Team" lose to named humans every time.

The operational detail people miss is the consistency: the same name, spelled the same way, linked to the same bio page, everywhere that person writes. The bio page should say who

the person is, what they have done, and link to their profiles elsewhere. This is how an engine assembles a credible author entity rather than a scattering of bylines.

How to verify: click every byline on your site. Each should land on a real bio, and no page should say "Team."

Phase 2: The structural fixes (ship this week)

Five fixes that make your existing content machine-readable. This is where the "60% of the opportunity is structural" finding lives.

Fix 4: Install schema.org/Article on every blog post

Minimum fields: `headline`, `author`, `datePublished`, `dateModified`, `description`. Use JSON-LD, not microdata: it is one script block in the page head, generated from data you already have in frontmatter or your CMS.

How to verify: run every templated page type through [Google's Rich Results Test](#). Because Article markup is usually implemented once in a template, one page passing generally means the type passes, but check one post, one older post, and any page that renders differently.

Fix 5: Add schema.org/FAQPage to any page with Q&A content

This is the highest-leverage schema type for GEO, because it maps one-to-one to how generative engines pattern their answers: question in, direct answer out.

The operational move: literally take the questions people ask you in sales calls and put them on the relevant page with structured FAQPage markup. Not the questions you wish they asked. The real ones, phrased the way a customer phrases them, each answered in two or three plain sentences. The markup then mirrors exactly what is visible on the page: markup for content that is not on the page is a penalty pattern, not a shortcut.

How to verify: Rich Results Test again, plus a read-aloud check: if an engine quoted only one of your answers, would the questioner be satisfied?

Fix 6: Add schema.org/Organization to your root layout

`name`, `url`, `logo`, `sameAs` (your social profiles), `description`. This is how engines build a canonical identity for your brand: one entity, with confirmed profiles, rather than several similarly named maybes.

Use the same description here that you use on your profiles and in your llms.txt. Consistency across sources is itself the signal.

Fix 7: Rewrite your top 5 pages to lead with the answer

For each of your five most important pages, the first two paragraphs should directly answer the question the URL implies. The background, the story, and the framing come after. Not the other way around.

This inverts how most marketing pages are written, which is why it works: when an engine retrieves your page for a query, it is looking for the answer near the top. If your first three paragraphs are scene-setting, the engine quotes a competitor who got to the point.

The mechanical version: for each page, write down the question its URL implies. Then check whether paragraphs one and two answer it. If not, move the answer up and push the story down. You rarely need to write anything new.

Fix 8: Add a TL;DR block at the top of every long-form page

Label it clearly ("**TL;DR:**" or "Key takeaway:"). Two to three sentences max. The test: if an LLM were to quote only this block, the user's question should be meaningfully answered.

Write it last, after the page is done, and make it a real summary with the specifics in it, not a teaser that withholds the point to force a scroll. Engines extract; they do not tease.

Phase 3: The depth work (ship over 2–3 weeks)

Two fixes that build the topical authority signal, the thing engines use to decide who "owns" a topic.

Fix 9: Build one deep pillar page per core topic

Pick your three most important topics. For each, build a 2,500 to 4,000 word pillar page that:

- Leads with the answer (fix 7 applies here doubly)
- Has section headings that are themselves direct claims, so the page's argument is legible from the heading outline alone
- Includes specific numbers, dated examples, and named sources
- Links out to three to five supporting cluster pages on your site

This is the slowest item on the list and the one most tempting to fake with generated filler. Don't. Engines deprioritize content that reads as synthetic, and a thin 4,000-word page is worse than a sharp 2,500-word one.

Fix 10: Cross-link the cluster

Each pillar page links to three to five supporting pages. Each supporting page links back to the pillar. That bidirectional structure is the point: it is how you demonstrate topical authority rather than assert it.

The operational check: draw your cluster on paper. Every arrow should run both ways, every supporting page should be reachable from its pillar, and no cluster page should be an orphan reachable only from the sitemap.

Fix 11: Add a "Sources" or "References" section where appropriate

Even on marketing pages. A list of the primary data, studies, or internal case studies that inform your claims. Two effects: LLMs prefer defensible content, and it makes your page genuinely more persuasive to humans.

Keep it honest: cite the actual sources behind your claims, dated, and only sources you have read. A decorative reference list is worse than none.

Phase 4: The measurement discipline (ship ongoing)

Fix 12: Build your citation scoreboard

Pick 20 to 50 queries your ideal customer actually types. Run them through ChatGPT, Perplexity, Gemini, Google AI Overviews, and Claude. Record who gets cited. Repeat every two weeks.

Operationally: keep the query list fixed so periods are comparable, run the same queries in the same engines, and log three things per query: were you named, was the description accurate, and which competitors appeared. A spreadsheet is enough.

This is the only metric that tells you whether the first 11 fixes are working. Without it, you're guessing.

What you should NOT do

A short list of GEO "advice" that is either useless or harmful:

- **Keyword stuffing for AI.** Engines are trained to recognize and penalize this. Don't.
- **Generating AI content to win AI citations.** The irony doesn't redeem the strategy: engines deprioritize content that reads as synthetic.
- **Paying for "AI SEO" tools before the structural work is done.** The tools are downstream of the fundamentals, and the fundamentals are the twelve items above.
- **Promising specific citation outcomes to stakeholders.** Generative engines are stochastic. Directional improvement is the right target, and the scoreboard is how you show it.

The operating rhythm

- **Founder or operator:** ship fixes 1 through 6 this week. They are near-free and compound.
- **With a content team:** hand them fixes 7 through 11 with a two-week window.
- **Ongoing:** the quick wins are one-time, the structural work is a quarterly review, and the scoreboard runs monthly to fortnightly. AI engines change retrieval behavior often enough that a set-and-forget audit decays within a couple of quarters.

The meta-point

GEO isn't a separate discipline from being a useful, credible source. It is the machine-readable version of the same idea: be clear, be specific, be sourced, be consistent, be cited. The brands that win will be the ones that structure their existing credibility so the AI engines can actually find it.

NEXT STEP

Want to go further

Want to see where you stand right now? Try our free tools:

- **Automation Advisor** at naavimlabs.com/tools/automation-advisor
- **Conversion Autopsy** at naavimlabs.com/tools/conversion-autopsy

Want this audited and executed for your brand? We run measured GEO programs that track citation share across the major AI engines. **Reach out**. No pitch, we just like fixing this specific problem.