

Most sites are invisible to AI search.

We have done a lot of research and tried a lot of things.
Here is what stuck.

MAY 2026

The words and the visuals carry the same message.

When a hero copy change ships without updating `og:description`, the schema headline, and the alt text, the page says one thing to humans and a different thing to LLMs.

We moved hero v2, the engraved-plate aesthetic, motion polish, and a three-word PR through the same queue. One operator means the copy and design stay in step.

The bridge between what humans see and what machines cite.

If the visible byline says a human but the JSON-LD says `Organization`, the LLM picks one and the citation comes out muddled.

We moved blog Article `author` to a `Person` object with `sameAs` linking to a Wikidata draft. Patched per-post; the template still defaults to `Organization` and that fix is the next ticket.

Crawlers learn your site from files no human reads.

`llms.txt`, `robots.txt`, sitemap, Link headers, prerendered HTML. If they disagree, crawlers index something other than what you meant.

We added an `llms.txt` index, prerendered seven routes, and wired Link plus Content-Signal headers. ArchConvert has an explicit AI-crawler allow-list; FluxCo uses a wildcard allow plus the Content-Signal posture.

Without it, you are guessing whether any of it worked.

Leading: which AI crawlers arrive on which routes. Lagging: AI referral, a person landing on the site because ChatGPT, Claude, or Perplexity sent them.

Cloudflare bot analytics gives us per-bot, per-route logs; we are still building the dashboard, so it is a manual scan right now. AI referral is the outcome the whole exercise is optimizing toward.

Technical side, figured out. Brand and messaging, still hopeless.

We are after feedback and a nudge in the right direction. Point us at what we should read, who we should follow, or what we have wrong.

Drop a comment on the long-form post, or reach out directly. Both channels are open.

Where the standards live.

- **llms.txt** · Jeremy Howard, Answer.AI, Sept 2024 · llmstxt.org
- **Schema.org** Article + Person + sameAs · W3C, Google Rich Results docs
- **E-E-A-T** · Google Search Quality Rater Guidelines
- **robots.txt** RFC 9309 · **Link headers** RFC 8288 · IETF
- **GEO** · Chen et al., Sept 2025 · arxiv.org/abs/2509.08919
- **AI citation patterns** · Yext AI Visibility Report, Oct 2025 · yext.com/blog
- **Cloudflare bot analytics** + Content Signals Policy · Cloudflare blog

Long version on the blog.

fluxcotech.com/blog/what-i-learned-shipping-30-marketing-prs

VANCOUVER / AUCKLAND