

# Your 90-day AEO action plan

---

Spencer Goldade

## Your 90-day AEO action plan

*From Answer Engine Optimization: The Complete Practitioner's Guide*

---

A week-by-week execution plan for the Access, Structure, Authority framework. Weeks 1 through 4 cover technical access. Month 2 covers content restructuring. Month 3 covers entity and platform work. Print it. Put it on a wall. Check things off.

This PDF matches `90-day-plan.md` in the bundle. If you prefer a tracker, open `90-day-tracker.csv` in the same folder.

**Spencer Goldade**

[spencergoldade.ca/aeo](http://spencergoldade.ca/aeo)

# Your 90-day AEO action plan

---

You've read the theory. You understand how AI answer engines select and cite content. You know the Access, Structure, Authority framework. Now you need to implement it.

For illustration, picture a mid-market B2B SaaS company with a marketing site on Cloudflare, a blog on a headless CMS, and one dedicated developer day per week. The calendar below still applies; resize the hours to your team.

This is the execution plan for all three layers, broken into a week-by-week schedule organised by who does what. Three roles drive the work: your developer (or technical team), your content team, and your marketing or SEO lead. The sequence follows the framework: **Access first** (weeks 1–4), then **Structure** (month 2), then **Authority** (month 3). This order isn't arbitrary. Each layer depends on the one before it. Print this. Put it on the wall. Check things off.

---

## Weeks 1–2: technical foundation (the fast wins)

These two weeks belong to your developer or technical team. The work is straightforward configuration — no content creation, no committee approvals. Changes show results in 1 to 4 weeks because they remove barriers that prevent AI bots from accessing your site.

**Owner:** Developer / Technical Lead

### Day 1–2: Audit and fix robots.txt

Many sites unknowingly block AI crawlers through overly broad disallow rules or inherited configurations. You need explicit entries for every major AI bot.

1. Open `robots.txt` at `yourdomain.com/robots.txt`.
2. Check for blanket `Disallow: /` rules applied to all user agents.
3. Add explicit allow rules for: GPTBot, OAI-SearchBot, ChatGPT-User, ClaudeBot, PerplexityBot, Google-Extended, Applebot-Extended.
4. If you use a CMS like WordPress, check whether a plugin is overriding manual `robots.txt` edits.
5. Validate with Google's `robots.txt` tester in Search Console.

6. Deploy and verify by checking server logs for AI bot requests within 48 hours.

This is the single highest-priority technical task. If AI bots can't crawl your site, nothing else in this plan matters. Full `robots.txt` templates are in `robots-txt/` in this bundle.

## Day 2–3: Disable Cloudflare Bot Fight Mode

Cloudflare Bot Fight Mode (and similar heuristics under Super Bot Fight Mode) can silently block AI crawlers when security rules treat them like abusive scrapers.

1. Log in to your Cloudflare dashboard.
2. Navigate to Security > Bots.
3. Check whether Bot Fight Mode or Super Bot Fight Mode is enabled.
4. Disable it, or configure custom rules that allow known AI crawler IPs.
5. Test by running `curl -A "GPTBot" https://yourdomain.com/` from the command line and checking for a 200 response.
6. Monitor server logs for the next week to confirm AI bots are getting through.

If you can't fully disable Bot Fight Mode, create WAF custom rules that allowlist the specific IP ranges published by OpenAI, Anthropic, and Perplexity.

## Day 3–5: Create your `llms.txt` file

The `llms.txt` file gives AI systems a structured summary of your site. Early adoption costs almost nothing and signals that you understand how AI systems consume content.

1. Create a plain text file at `yourdomain.com/llms.txt`.
2. Include: a one-paragraph site description, your primary topic areas, key pages with brief annotations, and any content licensing terms.
3. Optionally create `llms-full.txt` with more detailed content for systems that support it.
4. Add both files to your sitemap.
5. Verify the files are accessible and not blocked by `robots.txt` or Cloudflare rules.

Time investment: 1–2 hours. Full template is at `llms-txt/llms-sample.txt` in this bundle.

## Day 5–7: JavaScript rendering audit

AI crawlers handle JavaScript poorly. Some don't execute it at all.

1. Use Google's Rich Results Test or a tool like Rendertron to see what your pages look like without JavaScript.

2. Check your top 20 pages. If content is missing without JS, you have a rendering problem.
3. Prioritise server-side rendering (SSR) or static generation for pages you want AI systems to index.
4. As a temporary fix, ensure critical content appears in the initial HTML response before any JavaScript executes.
5. Test key pages with `curl` to see the raw HTML response.

Focus on your most important pages: homepage, pricing, top blog posts, and product pages.

## Day 7–10: Page speed and crawl budget

AI bots have crawl budgets just like Googlebot. If your pages are slow, bots crawl fewer of them.

1. Run Lighthouse or PageSpeed Insights on your top 20 pages.
2. Target a Time to First Byte (TTFB) under ~400ms for key pages as a planning goal. Use Core Web Vitals reports to sanity-check against real-user data.
3. Fix the obvious wins: compress images, enable browser caching, minimise render-blocking CSS and JavaScript.
4. Check server logs to see how many pages AI bots crawl per session. If the number is low (under 50 pages per visit), speed is likely a factor.

**Week 1–2 Checkpoint:** By day 10, your developer should have completed all five tasks. Verify by checking server logs for AI bot traffic. If you see GPTBot and ClaudeBot making successful 200 requests to your key pages, the technical foundation is in place.

---

## Weeks 3–4: schema implementation

Schema markup tells AI systems what your content means, not just what it says. A page with proper schema is easier for AI to parse, classify, and cite. This phase requires developer and content collaboration.

**Owner:** Developer + Content (collaboration)

### Week 3: Organization and FAQPage schema

Start with these two schema types — most immediate AEO value for the least complexity.

**Organization schema:**

1. Create a single Organization JSON-LD block for your homepage.
2. Include: `name`, `url`, `logo`, `sameAs` (all official social profile URLs), `contactPoint`, `foundingDate`, `description`.
3. Verify every field matches your Google Business Profile and Wikipedia entry (if you have one).
4. Validate with Google's Rich Results Test and Schema.org's validator.
5. Deploy to homepage and About page.

**FAQPage schema:**

1. Identify all pages with Q&A-style content.
2. Add FAQPage JSON-LD to each page, wrapping each question–answer pair.
3. Make sure the visible page content matches the schema content exactly. Google penalises mismatches.
4. Validate each page.

Templates live in `schema/` in this bundle.

**Week 4: Article, HowTo, and Person schema**

Once Organization and FAQPage are deployed, add schemas for your content and authors.

1. Add Article schema to all blog posts and editorial content.
2. Add HowTo schema to any page with numbered step-by-step instructions.
3. Add Person schema for each author, linking to their author page, LinkedIn, and other professional profiles.
4. Cross-reference Person schema with Organization schema using `employee` or `member`.
5. Validate everything. Fix errors before moving on.

**Week 3–4 Checkpoint:** Run Google's Rich Results Test on 10 representative pages. Every page returns valid structured data with zero errors. Warnings are acceptable at this stage. Errors are not.

---

## Month 2: content restructuring

The technical foundation is in place. AI bots can reach your site, read your content, and understand your schema. Now restructure the content so AI systems can extract the best answers from it.

**Owner:** Content Team

### Weeks 5–6: BLUF rewrites on top pages

BLUF stands for Bottom Line Up Front. Put the answer first, then provide supporting details. AI systems strongly favour this pattern.

1. Pull your top 20 pages by organic traffic from GA4.
2. For each page, identify the core question the page answers.
3. Rewrite the opening paragraph to deliver the answer in the first 2–3 sentences.
4. Move background, context, and methodology below the answer.
5. Keep each rewritten opening under 50 words.

Don't rewrite your entire site. Start with the top 20. Measure the impact. Then expand.

### Weeks 6–7: Heading overhaul to question format

AI systems parse headings to understand page structure. H2 headings written as questions align directly with how users prompt AI systems.

1. Audit H2 headings on your top 30 pages.
2. Rewrite label-style headings as questions or complete statements:
  - “Pricing” → “How Much Does [Product] Cost?”
  - “Features” → “What Features Does [Product] Include?”
  - “Implementation” → “How Do You Implement [Product]?”
3. Ensure each H2 section contains a direct answer to its heading question in the first sentence.
4. Aim for at least 60% of H2s in question format.

Check your Google Search Console query data for what people actually ask.

## Weeks 7–8: Content cluster gap analysis

AI systems favour authoritative content hubs over isolated pages.

1. List your primary topic areas (5–10 core topics).
2. For each topic, map the existing content you have (pillar page, supporting articles, FAQs).
3. Identify gaps where competitors have content you don't.
4. Prioritise 3–5 new pages that would complete your most important content clusters.
5. Create a content calendar for producing those pages over the next 30 days.

Use Semrush's Topic Research or Ahrefs' Content Gap analysis. Focus where AI systems are already citing competitors.

**Month 2 Checkpoint:** Top 20 pages have BLUF-structured openings. At least 60% of H2 headings on those pages are question-format. You have a documented content gap analysis with 3–5 priority pages identified.

---

## Month 3: entity and platform work

AI systems don't just evaluate your website. They evaluate your brand's presence across the entire web.

**Owner:** Marketing / SEO Lead

### Week 9–10: Knowledge Panel and entity consistency

1. Search your brand name on Google. If a Knowledge Panel appears, verify every detail.
2. If no Panel exists, claim your Google Business Profile and complete your Wikidata entry.
3. Audit entity consistency across your website, Google Business Profile, LinkedIn, Crunchbase, and directories.
4. Fix inconsistencies — even “Inc.” vs “Incorporated” can cause AI systems to see two different entities.
5. If you have a Wikipedia article, review it for accuracy. If not, check whether you meet notability criteria.

### Week 10–11: LinkedIn and Reddit presence

AI systems trained on web data pull heavily from LinkedIn and Reddit. These platforms appear in AI training data at disproportionate rates.

**LinkedIn:**

1. Ensure your company page has a complete, keyword-rich description.
2. Verify key team members have active profiles listing your company.
3. Publish 2–3 thought leadership posts per week from leadership on topics related to your core expertise.
4. Engage in relevant LinkedIn groups and comment threads.

**Reddit:**

1. Identify subreddits relevant to your industry.
2. Contribute genuinely useful answers (not marketing content) in threads about your topic area.
3. If someone asks about tools or companies in your space, it's fine to mention your company, but only alongside genuinely helpful context.
4. Never astroturf. Reddit communities detect promotional accounts quickly.

**Week 11–12: Wikipedia assessment and third-party mentions**

Wikipedia is the single most-cited source in AI training data.

1. If you have a Wikipedia article: read it and flag inaccuracies. Request corrections through Wikipedia's edit process. Don't edit your own article (conflict of interest).
2. If you don't have one: check whether you meet notability guidelines.
3. Identify 5–10 industry publications or review sites where your brand should appear but doesn't.
4. Pitch guest articles, seek product reviews, or submit to relevant directories.
5. Each third-party mention becomes a data point AI systems use when deciding whether to cite you.

**Month 3 Entity Checkpoint:** Knowledge Panel claimed and accurate. Entity data consistent across major platforms. LinkedIn and Reddit activity underway. List of third-party mention targets with outreach in progress.

**Month 3 ongoing: measurement setup**

Without measurement, you're guessing.

**Owner:** SEO Lead / Analytics Team

## **GA4 channel configuration**

1. Create the **AI Search** channel group in GA4 with regex rules for ChatGPT, Perplexity, Gemini, Copilot, and Claude referral domains.
2. Set it as your primary reporting view.
3. Build a simple dashboard showing AI referral sessions, conversion rate, and top landing pages.
4. Compare AI referral conversion rates against organic search conversion rates.

Exact patterns and regex strings: `measurement/ga4-ai-channel-setup.md` in this bundle.

## **Server log monitoring**

1. Use the `ai-bot-log-analysis.sh` script in `measurement/` in this bundle.
2. Schedule it to run monthly.
3. Flag any pages returning 403 or 429 responses to AI bots.
4. Track which pages get the most AI bot attention.

## **Manual testing baseline**

1. Build a prompt bank of 25–50 queries (brand, category, technical).
2. Run every prompt through ChatGPT, Perplexity, Gemini, and Copilot.
3. Record whether your brand appears, what position it's in, and which competitors show up.
4. Save this as your month-3 baseline.
5. Schedule monthly retesting using the same prompt bank.

Starter prompts: `measurement/prompt-bank-starter.md` in this bundle.

## **First citation tracking report**

1. Record share of voice (% of test prompts where your brand appears).
2. Record AI referral sessions and conversion rates from GA4.
3. Record AI bot crawl volume and top pages from server logs.
4. Record branded search trends from Google Search Console.
5. Note factual errors — these become entity optimisation priorities.

Scorecard template: `measurement/monthly-scorecard.csv` in this bundle.

---

## The 90-day milestone: what “working” looks like

After 90 days, don't expect a revolution. AEO is a compound strategy.

### Signs your AEO work is gaining traction:

- Server logs show GPTBot, ClaudeBot, and PerplexityBot returning 200 responses.
- Schema validates cleanly — zero errors across all types.
- Your brand appears in some AI responses for target queries. Even 10–20% is meaningful at this stage.
- AI referral traffic in GA4 is present and trackable, even if absolute numbers are small.
- Entity data is consistent across platforms.
- Server logs show AI bots crawling deeper into your site compared to your pre-AEO baseline.

### Signs something needs attention:

- AI bots still getting 403 or 429 errors.
- Schema validation shows errors (not just warnings).
- Your brand never appears in AI responses.
- AI systems describe your brand inaccurately.
- AI referral conversion rates are low.

### What to do at day 91:

Review your monthly scorecard. Identify which measurement signals are strongest and which are weakest. The weak signals point to your next 90-day focus.

AEO is not a project with a finish date. It's a practice that becomes part of how you publish content, structure data, and build brand presence. The 90-day plan gets you started.