

CHATGPT & AI VISIBILITY

Get named when buyers ask ChatGPT or Perplexity who to pick

Why recommendations go to competitors today — and the exact crawler rules, copy patterns, schema graph, and 10-minute monthly check that put your brand on the shortlist.

KEY BENCHMARKS INSIDE THIS DOSSIER

30 Checks
Simple 6-pillar audit covering crawlers, readability, trust & citation clarity

92 / 100
Verified score live on cr8vstacks.com across 8 major AI search crawlers

< 1.2s Read
Delivered in clean server HTML so AI fetchers see every word immediately

1 Clean Graph
Unified company, service, and FAQ identity with zero plugin conflicts

INSIDE THIS EDITION

Six chapters

01 **Why buyers ask ChatGPT before they click Google**
A quiet shift in B2B discovery means buyers shortlist 2–3 trusted options inside AI chat before opening a single browser tab.

02 **How bots read raw HTML & Cloudflare firewall rules**
Separate live citation indexers (`OAI-SearchBot` , `PerplexityBot` , `Claude-SearchBot`) from offline training scrapers cleanly without WAF blocks.

03 **Publishing /llms.txt cleanly without menu bloat**
Give coding agents and retrieval bots a crisp `< 2 KB` Markdown reading guide under the `llmstxt.org` specification.

04 **Answer-first pages, HTML tables & unified @graph schema**
Replace poetic slogans with 25–40 word answer capsules, native HTML `<table>` pricing specs & connected company JSON-LD.

05 **The 10-minute monthly check inside Free AEO Checker**
Spot caching regressions, accidental `403` blocks, or disconnected schema nodes monthly before rankings slip.

06 **Your 4-week implementation roadmap & impact benchmarks**
Sequence crawler access, `/llms.txt` , answer-first rewrites & prompt testing across four crisp weekly passes.

CHAPTER 01

Why buyers ask ChatGPT before they click Google

Five years ago, a buyer looking for a custom website or software vendor skimmed ten links on Google. Today that initial shortlist happens inside ChatGPT Search, Perplexity, or Claude in under thirty seconds.

3 in 5 B2B buyers	Compare vendors inside AI assistants before shortlisting agencies to contact directly.
< 1.2 sec budget	Live AI fetchers give your server response before dropping slow scripts and moving to a clearer site.
30-point audit	Every factor determining whether your brand is cited or skipped fits cleanly into six measurable pillars.

Across the sites we audit inside Free AEO Checker, the gap is rarely company quality—it is basic readability:

92 / 100 cr8vstacks.com verified live grade	**8 / 8** AI search bots returning HTTP 200
42% Sites silently blocking PerplexityBot	**0.9s** Server-rendered HTML delivery

What changed between classic SEO and AI recommendations

Google rewarded keyword density, backlink volume, and ten links per page. Answer engines compress twenty candidate tabs into two named recommendations and expect structured proof right inside initial HTML.

Real takeaway: A pretty homepage built only for human scrolling leaves ChatGPT Search with nothing concrete to verify or quote when your next buyer asks who to hire.

Ten Blue Links Era (2015–2023)

Buyers clicked 5–8 tabs, scanned vague hero copy, opened pricing pages, and compared agencies manually.

AI Shortlist Era (2025+)

Models fetch raw HTML once, extract exact capabilities, price ranges & case proof, and name 2–3 trusted options inline.

KEY RULE

Being readable only to human browsers is no longer enough: if AI retrieval agents cannot extract your capabilities inside raw HTML, you never enter the shortlist.

CHAPTER 02

How bots read raw HTML & Cloudflare firewall rules

Unlike classic two-wave Googlebot indexing, live AI fetchers parse your initial HTTP 200 payload directly in milliseconds—and frequently get dropped silently by Cloudflare Bot Fight Mode or JS CAPTCHAs.

Blocking `GPTBot` in `/robots.txt` protects content from offline training **without** hurting ChatGPT Search citations—as long as `OAI-SearchBot` and `ChatGPT-User` stay explicitly allowed.

Bot User-Agent	Vendor & Role	Trains AI?	Recommended Rule
OAI-SearchBot	OpenAI · ChatGPT Search index	No	PASS
ChatGPT-User	OpenAI · Live buyer chat fetch	No	PASS
PerplexityBot	Perplexity · Live citation index	No	PASS
Claude-SearchBot	Anthropic · Claude web search	No	PASS
GPTBot	OpenAI · Foundation model training	Yes	Policy choice
ClaudeBot	Anthropic · Model training corpus	Yes	Policy choice

The hidden Cloudflare Bot Fight Mode trap

Roughly two out of five sites we audit already allow search bots inside `/robots.txt` —only for Cloudflare WAF or security plugins to silently serve an HTTP `403` JavaScript challenge before a single word loads.

public_html/robots.txt — Clean Search Bot Separation

```
User-agent: OAI-SearchBot
Allow: /

User-agent: PerplexityBot
Allow: /

User-agent: Claude-SearchBot
Allow: /

Disallow: /wp-admin/
Sitemap: https://yourdomain.com/sitemap_index.xml
```

Inside Cloudflare Security -> WAF Custom Rules, skip Bot Fight / JS Challenge specifically when `cf.verified_bot_category` matches Search Crawlers—never leave `/wp-admin/` exposed.

KEY RULE

Confirm `HTTP 200` SSR text via `curl -A "OAI-SearchBot"`, explicitly allow search indexers inside `/robots.txt`, and allowlist verified crawler IPs in your WAF.

CHAPTER 03

Publishing /llms.txt cleanly without menu bloat

Proposed at `llmstxt.org`, `/llms.txt` hands AI assistants your canonical company summary, core services, case studies, and tools directly at domain root without navigation menus or cookie modals.

Myth: "`/llms.txt` ranks me #1 on ChatGPT"

False. Adding `/llms.txt` does not boost authority on its own or override blocked crawlers and weak service copy.

Reality: "A clean 2-KB reading map for agents"

True. Under `llmstxt.org`, it hands coding tools and retrieval agents your canonical URLs without nav clutter or menus.

STANDARD `LLMSTXT.ORG` 4-BLOCK ORDER

1. **# Official Brand Title (H1)**
   Exact registered business name matching your `Organization` schema.
2. **> Two-Sentence Blockquote**
   Plain-English explanation of what you build, who you serve, and offices.
3. **Core Services Links (`##`)**
   Annotated `[Title](https://...)` links to pricing, flagship builds, and guides.
4. **Proof & Tools Section (`##`)**
   Direct links to case studies, free calculators, and verified contact/legal pages.

A clean `/llms.txt` file takes fifteen minutes to ship

Keep yours strictly under 80 lines, plain UTF-8 Markdown, served at `HTTP 200` from your root domain. Link only evergreen canonical URLs.

Skip marketing fluff inside descriptions

Write clear 8–12 word descriptions after each colon (`: Fixed-scope WordPress builds from $3,500`) instead of adjectives.

Confirm `Content-Type: text/markdown; charset=utf-8`

Ensure your caching plugin does not wrap `/llms.txt` inside a themed HTML 404 or redirect chain.

cr8vstacks.com/llms.txt — Minimal Valid Reference

```
# Cr8v Stacks
> Custom WordPress, WooCommerce, and AI-ready web engineering studio.

## Core Services
- [Custom Websites](https://cr8vstacks.com/services/): Fixed-scope web builds.
- [Free AEO Checker](https://cr8vstacks.com/toolkits/free-aeo-checker/): Live 30-point audit.
- [Field Guides](https://cr8vstacks.com/guides/): Playbooks for founders and leads.
```

KEY RULE

Keep `/llms.txt` focused under 80 clean lines, served as UTF-8 Markdown, linking only high-signal pages.

CHAPTER 04

Answer-first pages, HTML tables & unified @graph schema

Retrieval models lift passage chunks anchored around clear headings and validate claims against connected `Organization`, `Service`, and off-site `sameAs` citations.

4 -> 1 JSON-LD	Replacing fragmented plugin snippets with one server-rendered `@graph` cuts ambiguity immediately.
5+ SameAs Links	Matching LinkedIn, GitHub, Clutch, Google Business Profile & domain records proves off-site authenticity.

One `@graph`, zero plugin clashes

Many WordPress sites emit 3 separate `Organization` islands from SEO, review, and local plugins. Merging under `@graph` creates one undeniable source of truth.

Deterministic `@id` anchors

Link `Service` -> `provider: {"@id": "/#organization"}` and `WebPage` -> `publisher` so models resolve your entity graph in a single pass.

Off-site provenance: models double-check claims outside your domain

Clean schema on your site only finishes half the job. Retrieval models cross-verify your claims against directories, legal policies, and public social profiles.

Consistent entity facts everywhere
Keep company name, core services, named leadership, and physical office addresses (for example Lagos & Ibadan) word-consistent across your footer, Privacy Policy, LinkedIn, and Google Business Profile.

Include verified review & case-study URLs inside `sameAs`
Explicitly connecting verified Clutch, GitHub, and YouTube handles helps answer engines tie third-party praise back to your domain.

wp_head — Connected @graph Reference

```
{
  "@context": "https://schema.org",
  "@graph": [
    { "@type": "Organization", "@id": "https://yourdomain.com/#organization", "sameAs": ["..."]
},
    { "@type": "Service", "provider": { "@id": "https://yourdomain.com/#organization" } }
  ]
}
```

KEY RULE
Output one server-rendered `@graph` cleanly in `<head>` and pair every key `H2` question with an unambiguous factual answer capsule.

CHAPTER 05

The 10-minute monthly check inside Free AEO Checker

Plugin updates and CDN rule edits quietly introduce broken crawler headers or JS-locked copy. A repeatable 10-minute monthly check protects citation momentum all year.

STEP 01

Run Free AEO Checker

Scan your homepage and top service page inside `/toolkits/free-aeo-checker/` after any plugin or template deploy.

STEP 02

Confirm `200 OK` on Key Bots

Verify `OAI-SearchBot` and `PerplexityBot` still see 100% SSR HTML with zero Cloudflare `403` blocks.

STEP 03

Test 5 Real Buyer Prompts

Ask ChatGPT Search and Perplexity real vendor-selection questions in your niche and check which URLs earn citations.

STEP 04

Refresh `/llms.txt` & Proof

Add newly published case studies, guides, and free calculators to `/llms.txt` cleanly once per month.

LIVE DIAGNOSTIC SCOREBOARD

cr8vstacks.com · Verified WordPress SSR

92 / 100 · GRADE A

8 / 8 Allowed	< 1.1s SSR	Unified @graph
AI SEARCH CRAWLERS	ZERO JS HYDRATION WAIT	ORGANIZATION + SERVICES + FAQ

Your repeatable AI visibility audit checklist

Use this grouped check before major launches and at the start of every month.

Access & Speed (Every Deployment)

- ○ HTTP 200 confirmed for `OAI-SearchBot`, `ChatGPT-User`, `PerplexityBot` & `Claude-SearchBot`
- ○ No accidental `X-Robots-Tag: noai` headers or WAF challenges on service & pricing pages
- ○ Critical hero copy, tables & FAQs visible cleanly with JavaScript disabled

Copy, Schema & Provenance (Monthly)

- ○ Every top service `H2` opens with a 25–40 word verbatim answer capsule
- ○ `Organization`, `Service`, and `FAQPage` linked inside one clean `@graph` with zero schema warnings
- ○ `/llms.txt` updated with new flagship work, case studies, and interactive utilities

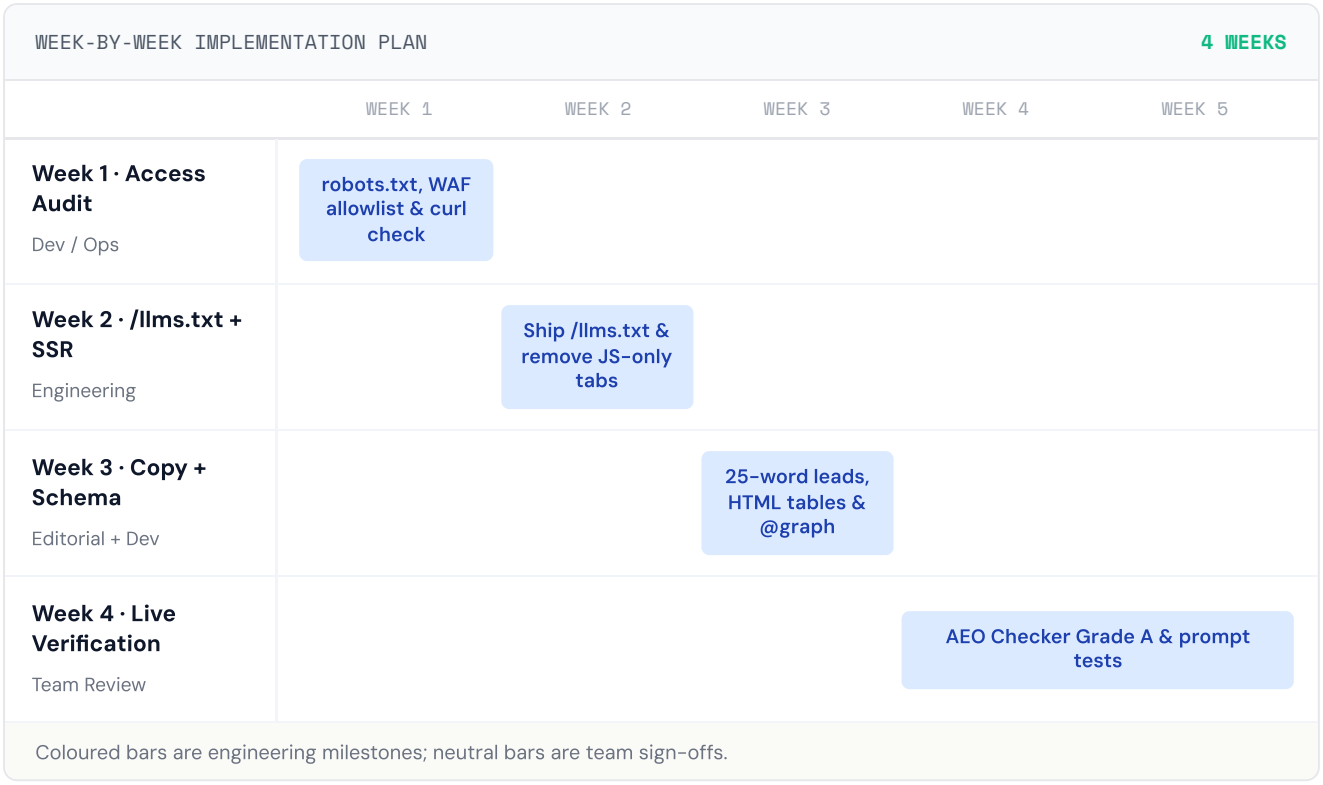
KEY RULE

Grade-A AI citations compound when technical crawlability and answer clarity stay verified every 30 days.

CHAPTER 06

Your 4-week implementation roadmap & impact benchmarks

You do not need a website rebuild to win AI visibility. Most teams ship the full stack cleanly in four weeks—starting with access checks in Week 1 and finishing with live prompt tests in Week 4.



FREE INTERACTIVE UTILITY

Free AEO Checker →

Run our live 30-point diagnostic across crawler rules, server HTML readability, `/llms.txt`, and `@graph` schema right now.

cr8vstacks.com/toolkits/free-aeo-checker

FREE INTERACTIVE UTILITY

Request AI Visibility Review →

Want our engineering team to audit your live stack and ship the fixes in one sprint? Book a 30-minute walkthrough.

cr8vstacks.com/discovery-call

Where four weeks of effort delivers the highest return

You get 80% of the citation lift from fixing crawl access, surfacing clear pricing/service tables in HTML, and unifying company identity.

Do not start with content volume alone. Fix crawler access and server readability first—then every clear page you publish enters the index immediately.

Impact on AI Shortlist Inclusion vs. Team Effort



KEY RULE

Always fix technical access (`robots.txt` , WAF, SSR HTML) before editing headlines so every improvement actually reaches crawlers.

RATHER TALK IT THROUGH?

Want an AI search audit of your own URL?

Bring your live domain to a 30-minute call with our team. We will run our live crawler and schema inspection on screen and show you the fastest wins.

Book a 30-minute call at cr8vstacks.com/discovery-call

Which stack fits
Whether to tune what you have, build natively, or migrate cleanly.

How long it takes
Realistic sprint weeks tailored to your pages, CMS, and integrations.

What it costs
Fixed-scope pricing confirmed before any engineering starts.

