

The AI visibility stack

Updated 9 September 2026 • Repository snapshot: 9 September 2026

Your next customer may be asking AI who to use. Start by measuring the answer, then work on the parts you can change.

This collection groups open-source and source-available tools by the job they do. Choose the part you need; installing the whole list is not a requirement for being named.

Repository stars, last-push dates and licence labels were checked on GitHub on 9 September 2026. They describe that snapshot, not a performance or security test. Hosting, model-provider fees and licence conditions may apply. Read the repository's licence before deploying it.

How to read this

Deploy: needs a server. **Run local:** a CLI, desktop app or build step. **Drop-in:** a library for an existing site. **Reference:** a collection to read.

Stars show interest. Judge quality yourself. “Check repository licence” means GitHub did not return one recognised licence label.

1. Measure it

Query the engines with buyer questions and read who they name. Answer monitors and technical crawlers do different jobs: use the former to measure answers and the latter to inspect your site.

RUN LOCAL

Auriti-Labs/geo-optimizer-skill

Audits and tracks whether ChatGPT, Perplexity, Gemini and Google AI Overviews cite your site. CLI, Python and MCP.

★ 781 Last push: 2026-09-05 MIT

RUN LOCAL

mverab/eGEOagents

Measures and optimises content for ChatGPT, Perplexity, Gemini and Claude. Ships on PyPI as egeo.

★ 181 Last push: 2026-09-01 MIT

RUN LOCAL

alexpospekhov/searchstack-aeo

One CLI that monitors visibility across Google AI Overviews, ChatGPT, Perplexity, Claude and Grok, and is happy on a cron.

★ 99 Last push: 2026-05-04 MIT

DEPLOY

promptclarity/promptclarity

A self-hosted Docker monitor of how AI mentions your brand across ChatGPT, Gemini and Perplexity, with the sources each engine cited.

★ 3 Last push: 2026-01-08 AGPL-3.0

DEPLOY

StJudeWasHere/seonaut

A self-hosted technical audit crawler for finding crawlability and on-page SEO issues.

★ 780 Last push: 2026-05-23 MIT

RUN LOCAL

adityaarsharma/librecrawl-technical-seo-audit-mcp

A technical crawler exposed as an MCP server, so your own agent can check crawler access and structured data directly.

★ 39 Last push: 2026-07-28 MIT

RUN LOCAL

zubair-trabzada/geo-seo-claude

A GEO-first audit skill for Claude Code: citability scoring, AI crawler analysis, schema markup and platform-specific checks in one run.

★ 10,490 Last push: 2026-09-09 MIT

REFERENCE

amplifying-ai/awesome-generative-engine-optimization

A curated list of GEO resources if you want to go wider.

★ 500 Last push: 2026-04-14 Check repository licence

REFERENCE

bmpi-dev/awesome-seo

A broader SEO research and traffic reading list for when you want the fundamentals underneath the AI layer.

★ 2,786 Last push: 2026-08-23 MIT

2. Make your site AI-readable

Make your identity and content explicit. An llms.txt is a proposed format for an AI-readable site guide; Schema.org JSON-LD describes the entities and content on a page. Neither guarantees a citation.

RUN LOCAL

AnswerDotAI/llms-txt

The canonical llms.txt spec from Answer.AI, plus a CLI that expands your file into a full context document for an LLM.

★ 2,606 Last push: 2026-09-04 Apache-2.0

DROP-IN

garmeeh/next-seo

Manages your meta tags and JSON-LD in Next.js, which are the signals an engine parses to work out what a page is.

★ 8,512 Last push: 2026-07-29 MIT

DROP-IN

google/schema-dts

TypeScript types for Schema.org JSON-LD. Check the resulting markup against the visible page and actual entity details.

★ 1,239 Last push: 2026-09-01 Apache-2.0

DROP-IN

4hse/astro-llms-txt

Generates llms.txt at build time for Astro sites, with no ongoing effort.

★ 23 Last push: 2026-04-21 Check repository licence

DEPLOY

apify/actor-llmstxt-generator

Crawls a URL and produces an llms.txt for it, for people who do not want to run anything locally.

★ 32 Last push: 2026-05-25 Apache-2.0

3. Be crawlable and renderable for agents

Some crawlers do not execute JavaScript. Check the HTML your server returns, then decide which crawlers you allow. Rendering and access controls are separate from whether an engine chooses to cite you.

RUN LOCAL

stereobooster/react-snap

Pre-renders a single-page app to static HTML at build time, so a non-JS crawler sees the content.

★ 5,116 Last push: 2026-02-02 MIT

DROP-IN

prerender/prerender-node

Express middleware that serves a rendered HTML snapshot to crawlers. Its recent commits specifically add AI-crawler user-agents.

★ 921 Last push: 2026-06-12 Check repository licence

DROP-IN

ai-robots-txt/ai.robots.txt

A maintained list of AI crawler user agents and blocking rules. Inspect it before use: copying the supplied rules can block the crawlers you wanted to allow.

★ 4,105 Last push: 2026-09-08 MIT

RUN LOCAL

unclecode/crawl4ai

A crawler and scraper for extracting pages as text for your own agent workflows. Its output is not a reproduction of every AI engine's crawl.

★ 82,021 Last push: 2026-09-08 Apache-2.0

DEPLOY

firecrawl/firecrawl

A context API to search and scrape the web at scale. Useful for turning pages into the clean structured text agents ingest, and for watching how your content reads once stripped to text.

★ 178,154 Last push: 2026-09-09 AGPL-3.0

4. Publish citable content

Publish clear, sourced content in HTML that can be fetched. Static generators and content management systems give you different ways to maintain it.

RUN LOCAL

gohugoio/hugo

A very fast static-site generator that outputs pure HTML with no JavaScript dependency.

★ 89,759 Last push: 2026-09-09 Apache-2.0

RUN LOCAL

withastro/astro

A content-first framework that ships near-zero-JavaScript HTML by default, with content collections built in.

★ 62,422 Last push: 2026-09-09 Check repository licence

RUN LOCAL

jackyzha0/quartz

Turns Markdown or Obsidian notes into an interlinked static site.

★ 13,193 Last push: 2026-09-06 MIT

DEPLOY

TryGhost/Ghost

A publishing platform for articles and newsletters, with server-rendered pages and structured data.

★ 55,234 Last push: 2026-09-09 MIT

DEPLOY

payloadcms/payload

A headless CMS built on Next.js, so your content and your fast server-rendered pages live in one place.

★ 44,637 Last push: 2026-09-08 MIT

DEPLOY

strapi/strapi

A headless CMS, API-first, for rendering structured content into clean pages.

★ 73,099 Last push: 2026-09-09 Check repository licence

DEPLOY

directus/directus

Turns a database into a content and data platform. Source-available under custom licence terms; check whether your use needs a paid licence.

★ 37,822 Last push: 2026-09-09 Check repository licence

5. Become agent-usable

MCP lets you expose tools and data to an AI client. Choose what it may read or change, and add authentication for private operations.

RUN LOCAL

modelcontextprotocol/servers

The official reference collection of MCP servers you can copy as a starting point.

★ 90,180 Last push: 2026-09-03 Check repository licence

RUN LOCAL

modelcontextprotocol/python-sdk

The official Python SDK for building MCP servers and clients.

★ 24,246 Last push: 2026-09-08 MIT

RUN LOCAL

modelcontextprotocol/typescript-sdk

The official TypeScript SDK, for teams already on Node or serverless.

★ 13,355 Last push: 2026-09-08 Check repository licence

RUN LOCAL

jlowin/fastmcp

A fast, Pythonic framework for MCP servers with remote deploy and auth built in.

★ 27,577 Last push: 2026-09-09 Apache-2.0

DEPLOY

cloudflare/mcp-server-cloudflare

Cloudflare’s own remote MCP servers on Workers, a production template for an endpoint agents can reach over the network.

★ 4,172 Last push: 2026-09-01 Apache-2.0

DEPLOY

github/github-mcp-server

GitHub’s official MCP server, worth reading as a reference build for how a serious product exposes itself to agents.

★ 32,816 Last push: 2026-09-08 MIT

DEPLOY

aiptoeosis-labs/aci

An open tool-calling platform that hooks 600+ tools into an agent through one unified MCP server, if you would rather connect than build.

★ 4,889 Last push: 2026-05-28 Apache-2.0

REFERENCE

punkpeye/awesome-mcp-servers

A large catalog of existing MCP servers. Check it before you build one, because someone may have already built it.

★ 94,677 Last push: 2026-09-08 MIT

6. Monitor mentions and citations

Answer monitors show sampled mentions and citations. Analytics show visits, while server logs can show crawler requests. A crawler visit is not evidence of a recommendation.

RUN LOCAL

allinurl/goaccess

A real-time log analyser, and a direct way to count GPTBot, ClaudeBot and PerplexityBot in your own server logs.

★ 20,916 Last push: 2026-09-03 MIT

DEPLOY

matomo-org/matomo

Self-hosted analytics whose log-import feature lets you build AI-bot segments, which the beacon-only tools cannot.

★ 21,849 Last push: 2026-09-09 GPL-3.0

DEPLOY

elmohq/elmo

A self-hostable platform that tracks how the engines mention and cite your brand, with every metric auditable in the code.

★ 313 Last push: 2026-09-08 MIT

DEPLOY

ai-search-guru/getcito

A self-hosted dashboard that tracks your visibility over time across ChatGPT, Perplexity, Claude and Google AI Overview.

★ 409 Last push: 2026-09-07 Check repository licence

DEPLOY

letterstory/lettertrace

Open-source, bring-your-own-key answer-engine telemetry. You supply the model keys and it records how the engines answer for you over time, which keeps the running cost yours and the data in your hands.

★ 69 Last push: 2026-09-08 MIT

7. Own your entities

Keep the names and identifiers for your organisation, people and products consistent. These general Wikidata and Schema.org tools help manage that information.

RUN LOCAL

OpenRefine/OpenRefine

Cleans and reconciles organisation, people and product data against services such as Wikidata.

★ 11,990 Last push: 2026-09-09 BSD-3-Clause

DROP-IN

google/schema-dts

TypeScript types for Schema.org JSON-LD. Check the resulting markup against the visible page and actual entity details.

★ 1,239 Last push: 2026-09-01 Apache-2.0

DROP-IN

spencermountain/wtf_wikipedia

Parses Wikipedia markup so you can pull and watch your own entity data instead of finding out late that it drifted.

★ 856 Last push: 2026-09-04 MIT

Being found is the goal.

Found by AI asks seven AI engines your buyers' questions and keeps every answer word for word. Where an engine exposes its web searches, we keep those alongside the answer. Your AI can read your measurements over MCP.

The readiness method and audit suite are open source. The hosted visibility measurement is a service we run. Stay Found re-measures weekly and turns the results into a prioritised fix plan.

Start with the free scan: ChatGPT and Gemini, about a minute, no signup or card. Read the answers before you choose what to do next.

[Read the readiness method and audit suite](#) · [How the measurement works](#)

Scan your site free →

Tech Horizon Labs · areyoufoundbyai.com · [OpenSEO and Found by AI](#) · [Build your own website](#)