APPS BY VIKAS

Meet Vikas



GenzLudo
Ludo with 3D characters and up to 8 players. A first-of-its-kind no-teaming-up feature keeps play fair. Try it - you'll love the difference.
Download on Google Play →



Kachu
Too much scrolling straining your neck? Add your addictive apps and complete daily neck exercises to unlock them.
Download on Google Play →

COMPLETE STRATEGY & TECHNICAL GUIDE · @VIKASBUILDSAPP

Beyond Traditional SEO: How to Get Website Traffic from AI Agents (GEO)

A plain-English blueprint explaining why traditional SEO alone is no longer enough, how AI agents discover and read websites, and the free tools you can use to make your site AI-ready.

Curated by **Vikas Yadav** · VikasBuilds (vikasbuilds.com)

1. Why Traditional SEO Alone No Longer Brings Traffic

For years, getting website traffic followed one simple formula: pick keywords, write blog posts, build backlinks, and wait to rank on Google's first page. Users would type a query, open 5 to 10 blue links, and compare websites themselves.

That user behavior has fundamentally changed. Today, millions of people skip traditional search results and ask AI tools directly—such as **ChatGPT, Perplexity, Google AI Overviews, Claude Code, and Codex**. Instead of giving users 10 links to browse, the AI recommends 1 to 3 trusted websites and summarizes the answer. If AI systems cannot find or read your website, users never even know you exist.

Two Types of Website Visitors Now Exist

Every website on the internet now serves two completely different audiences at the same time:

Factor	Visitor Type 1: Basic Humans	Visitor Type 2: AI Agents
How They Browse	Open your site in Chrome, Safari, or mobile browsers and scroll visually.	Fetch your URL programmatically via bots or terminal tools (Claude Code, Codex, ChatGPT).
What They Care About	Clean UI design, colors, animations, layout, and visual trust signals.	Clean text, clear headings, structured facts, pricing, and fast machine-readable data.
Why They Leave	Slow loading, cluttered design, or confusing navigation.	Heavy JavaScript blocks, missing structure, blocked crawlers, or bloated HTML noise.
Optimization Method	UI/UX + Traditional SEO (Keywords, backlinks, page speed).	**GEO — Generative Engine Optimization** (Discovery + AI Readability).

> **Key Takeaway:** Traditional SEO is not dead, but it is no longer sufficient on its own. You now need **SEO + GEO (Generative Engine Optimization)** working together so both search engines and AI agents recommend your website.

2. The 2-Step GEO Loop: Discovery + Readability

Many website owners make the mistake of adding an AI text file to their site and expecting instant traffic. In reality, getting traffic from AI systems works in two strict stages—in this exact order:

Stage	Core Question	What Actually Happens
Step 1: Discovery	How does the AI system find your website?	AI systems do not crawl the entire internet blindly from zero. They discover URLs through search indexes (Google, Bing, and their own search crawlers), public brand mentions, and direct links.
Step 2: Readability	Once it visits, can it read your content accurately?	The system inspects your page. Chat assistants read clean Server-Side Rendered (SSR) HTML, while coding agents prefer clean Markdown. Clear text and matching structured data let them cite you without guessing.
Step 3: Traffic	Why does the human visit your site?	The AI links to your page as a verified source inside its answer, sending high-intent human visitors straight to your website.

3. Step 1: Discovery (How AI Agents Find Your Website)

Before any AI assistant or coding agent can read your website, it must first know your URL exists. Here are the four verified channels through which AI systems discover websites:

A. Search Engine Indexes (Google, Bing & AI Search Crawlers)

Major AI answer engines combine their own crawlers with established search indexes:

- **Google AI Overviews & AI Mode:** Pull directly from Google's core Search index (crawled by `Googlebot`). Google uses a technique called *query fan-out*—running multiple related searches across indexed pages to assemble one cited answer.
- **ChatGPT Search:** Combines OpenAI's own search indexer (`OAI-SearchBot`) with third-party search providers, including Microsoft Bing, plus live page fetches (`ChatGPT-User`).
- **Perplexity & Claude Search:** Use their own dedicated search crawlers (`PerplexityBot` and `Claude-SearchBot`) alongside live user-triggered fetches to index and cite current pages.

> **Action Required:** Verify your domain and submit your `sitemap.xml` to both Google Search Console and Bing Webmaster Tools (and enable IndexNow on Bing for faster indexing).

B. Third-Party Brand Mentions (The Citation Loop)

When a user asks an AI for recommendations (such as *"best tools for..."* or *"reliable service providers for..."*), the AI evaluates public discussions across the web—including **Reddit threads, GitHub repositories, Quora answers, YouTube descriptions, comparison articles, and trusted directories**. When your brand name and website URL appear consistently in these places, AI systems discover your domain and visit your site to verify your claims.

C. Direct URLs in Developer & Research Workflows

Users routinely copy links from social bios, YouTube descriptions, or GitHub READMEs and paste them into Claude Code, Codex, Cursor, ChatGPT, or Perplexity with prompts like *"Check this website and summarize their pricing."* In these workflows, the AI fetches your URL directly on demand.

D. Internal Links & Clean Sitemaps

When a search crawler lands on your homepage via an external link, clear internal links and an up-to-date `sitemap.xml` ensure it discovers your pricing, documentation, product, and FAQ pages instead of stopping at the homepage.

4. Step 2: Readability (How AI Systems Read Your Website)

Once an AI crawler or agent reaches your URL, visual polish does not help it. Animations, complex CSS layouts, and client-side JavaScript that hides text until a browser executes scripts can prevent crawlers from seeing your real content.

Two Different AI Readers—One Website (No Cloaking)

A common myth is that you must build two completely separate websites—one for humans and a hidden one for bots. **Never show different facts, prices, or claims to bots than you show to humans.** Search engines classify deceptive differences as Cloaking and can penalize your site. Instead, serve the exact same information from one website in formats both types of AI readers can parse cleanly:

- **AI Chat & Search Assistants (ChatGPT Search, Perplexity, Google AI Overviews):** read standard **Server-Side Rendered (SSR) HTML**—the exact HTML a browser receives before heavy client-side JavaScript runs.
- **AI Coding Agents (Claude Code, Cursor, Codex, GitHub Copilot):** prefer clean **Markdown** when fetching URLs, typically requesting it via the HTTP header `Accept: text/markdown`.

The 4 Technical Layers of AI Readability

Layer 1: Clean Server-Side Rendered (SSR) HTML

This is the single most important technical foundation. Most AI search crawlers do not reliably execute heavy client-side JavaScript. Ensure your main headings, paragraphs, prices, tables, and links exist directly in the initial HTML response sent by your server (using static generation or SSR rather than a blank JavaScript shell).

Layer 2: HTTP Content Negotiation (`Accept: text/markdown` + `Vary: Accept`)

To serve coding agents cleanly without touching the human experience:

- When a human visits via Chrome or Safari, the browser sends `Accept: text/html` → your server returns the full visual website.
- When a coding agent requests your page with `Accept: text/markdown` → your server (or edge layer) returns clean Markdown containing the exact same page text, saving the agent's context tokens.
- **Critical Caching Rule:** Always include the response header `Vary: Accept` on negotiated routes. Without `Vary: Accept`, a CDN or browser cache could accidentally serve raw Markdown to a human visitor or heavy HTML to an AI agent.

Layer 3: Exact `robots.txt` Rules (Search vs. Live Fetch vs. Training)

Many guides mistakenly lump all AI bots together. In reality, major AI platforms separate **Search Indexing, Live User Fetches, and Model Training** into distinct user-agents—so you can allow AI Search citations even if you choose to block model training:

Platform	Search Indexing (Allow for Citations)	Live User Fetch (On-Demand Visits)	Model Training (Optional Choice)
OpenAI / ChatGPT	`OAI-SearchBot` Indexes pages for ChatGPT Search citations.	`ChatGPT-User` Fetches a URL when a user asks ChatGPT to read it.	`GPTBot` Crawls web data for OpenAI model training.
Anthropic / Claude	`Claude-SearchBot` Indexes pages for Claude search results.	`Claude-User` Retrieves pages in response to live user queries.	`ClaudeBot` Collects web content for Anthropic model training.
Perplexity	`PerplexityBot` Indexes sites for Perplexity search (not foundation training).	`Perplexity-User` Fetches pages during live user queries.	*Not applicable* (Perplexity states bot is for search, not training).
Google	`Googlebot` Powers Google Search, AI Overviews & AI Mode.	Standard Search indexing & Google fetchers.	`Google-Extended` Controls Gemini/Vertex AI training & grounding only.

Layer 4: Structured Data (JSON-LD Schema) & the Honest Role of `/llms.txt`

- **JSON-LD Schema:** Add valid Schema.org markup (`Organization`, `Service` or `Product`, and `FAQPage`) that strictly matches the visible text on your page. Google officially confirms there is no special AI-only markup required for AI Overviews, but accurate structured data reduces ambiguity across search engines and AI assistants.
- **The Truth About `/llms.txt`:** Proposed by Jeremy Howard (Answer.AI) in September 2024, `/llms.txt` is a plain Markdown index file placed at your root domain. Be realistic about what it does: **it is NOT a confirmed ranking factor for Google Search, ChatGPT Search, or Perplexity**—Google explicitly states Search does not use AI text files, and AI search bots overwhelmingly crawl HTML directly. However, `/llms.txt` is low-effort to publish and genuinely useful for IDE coding agents and developer workflows (like Cursor, Claude Code, and MCP tools) that look for a clean directory of your docs or offerings.

5. Verified Tools & Open-Source Stack

Below is the verified list of tools—including exact plan requirements, compatibility, and real install commands so nothing breaks when you try it:

Tool / Skill	Best Used For	Verified Setup, Compatibility & Cost
llms-visibility (Evil Martians)	Free, code-based `Accept: text/markdown` routing, SSR checks & `robots.txt` signals	**Cost:** 100% Free & Open Source (MIT). Works in Claude Code, Codex, Cursor, Gemini CLI, and GitHub Copilot. **Install:** `npx skills add evilmartians/agent-skills --skill llms-visibility -a claude-code` (swap `-a` for codex, cursor, etc.; add `-g` for global install). **Repo:** github.com/evilmartians/agent-skills
geo-seo-claude (Zubair Trabzada)	Deep GEO + SEO audit, AI crawler check, citability scoring & JSON-LD Schema	**Cost:** 100% Free & Open Source (MIT). Built specifically for Claude Code CLI (requires Python 3.10+ and Git). **Install (macOS/Linux):** `git clone https://github.com/zubair-trabzada/geo-seo-claude.git && cd geo-seo-claude && ./install.sh` (Windows: run `./install-win.sh` in Git Bash). **Run:** `/geo audit https://yoursite.com` **Repo:** github.com/zubair-trabzada/geo-seo-claude
Cloudflare: Markdown for Agents	Zero-code edge HTML-to-Markdown conversion (adds `x-markdown-tokens` header)	**Plan Requirement:** Included at no extra cost on Cloudflare Pro, Business, Enterprise, and SSL for SaaS plans (not available on the Cloudflare Free plan). **Location:** Cloudflare Dashboard → Agent Readiness > Markdown for Agents → Turn on *Serve Markdown to AI agents*. **Docs:** developers.cloudflare.com/.../markdown-for-agents
Unlighthouse (Bonus SEO Audit)	Site-wide technical SEO & Core Web Vitals scan	**Cost:** 100% Free & Open Source. **Run:** `npx unlighthouse --site https://yoursite.com` **Repo:** github.com/harlan-zw/unlighthouse
SerpBear (Bonus Rank Tracker)	Self-hosted daily Google keyword rank tracking	**Cost:** 100% Free & Open Source (self-hosted). **Repo:** github.com/towfiqi/serpbear

6. Copy-Paste Prompt for Claude Code or Codex

Open your website codebase in Claude Code or Codex and use this verified prompt to implement AI readability safely:

AUDIT & IMPLEMENT PROMPT FOR CLAUDE CODE / CODEX

```
Audit this website codebase and make it fully AI-agent and GEO ready without
breaking the human UI or causing search engine cloaking issues:

1. Server-Side HTML Check: Verify that all primary content (headings,
descriptions, pricing, FAQs, and links) renders directly in initial SSR or
static HTML without requiring client-side JavaScript execution.
2. robots.txt: Explicitly allow search and live-fetch bots — Googlebot, OAI-
SearchBot, ChatGPT-User, Claude-SearchBot, Claude-User, PerplexityBot, and
Perplexity-User. Ask me whether I want to allow or disallow training crawlers
(GPTBot, ClaudeBot, Google-Extended).
3. Content Negotiation: Add lightweight middleware that detects `Accept:
text/markdown` requests and returns clean Markdown of the exact same visible
page content. Include `Content-Type: text/markdown; charset=utf-8` and `Vary:
Accept` response headers so CDNs and browsers never mix up cached HTML and
Markdown.
4. JSON-LD Schema: Add valid Schema.org JSON-LD markup (Organization,
Service/Product, and FAQPage where applicable) that matches visible page text
verbatim.
5. /llms.txt (Bonus for Coding Agents): Generate a concise, accurate /llms.txt
index at the root domain linking to the site's key pages and core offerings.
```

Implementation Note: After running this prompt, test your endpoints with `curl -H "Accept: text/markdown" https://yoursite.com` to confirm clean Markdown is returned, and verify with standard browser navigation to confirm the human UI is completely unaffected.

7. Step-by-Step Action Checklist

Execute these steps in order so Discovery and Readability work together as one complete system:

Step	Pillar	Verified Action Item
Step 1	Discovery	Verify your site and submit `sitemap.xml` in Google Search Console (covers Google Search, AI Overviews, and AI Mode).
Step 2	Discovery	Verify your site, submit `sitemap.xml`, and enable IndexNow in Bing Webmaster Tools (supports Bing & ChatGPT Search discovery).
Step 3	Discovery	Place your consistent brand description and URL on GitHub, Reddit, YouTube, and relevant public profiles or directories.
Step 4	Readability	Confirm core page text renders in raw HTML (SSR/static) and check for crawl/speed errors using `npx unlighthouse`.
Step 5	Readability	Configure `robots.txt` to allow search/user bots (`Googlebot` , `OAI-SearchBot` , `ChatGPT-User` , `Claude-SearchBot` , `Claude-User` , `PerplexityBot`) and decide separately on training bots (`GPTBot` , `ClaudeBot` , `Google-Extended`).
Step 6	Readability	Enable `Accept: text/markdown` routing with `Vary: Accept` via `llms-visibility` (free) or Cloudflare Markdown for Agents (Pro+ plans), plus add `/llms.txt` for coding agents.
Step 7	Readability	Add matching JSON-LD Schema, run `/geo audit` via `geo-seo-claude` in Claude Code, and test how ChatGPT, Perplexity, and Claude cite your URL.

Built by VikasBuilds · vikasbuilds.com · @vikasbuildsapp

Need custom mobile app development, high-converting websites, AI agent readiness, or app promotions? Visit **vikasbuilds.com** or connect on Instagram at **@vikasbuildsapp**.

More things worth exploring

Find more free guides, checklists, and developer tools at vikasbuilds.com/resources →