

Sloka Agent

Framework, safety, revenue, and site positioning

This note is written for the owner of genaimakers.com. It is grounded in the site as it actually ships: a static HTML factory on Hostinger, PHP API proxies, Groq for generation, and Anthropic already used on House & Land AI. It is not a proposal to rebuild the stack in Next.js or LangChain.

Ship shape. Sloka Agent is live at genaimakers.com/sloka-agent.html (also /sloka; old /bhajan and bhajan-agent.html still resolve). It sits with House & Land AI and Trolley as a third proof-of-factory product — not a fake suite name, and not a new top-nav item. Users can request slokas, Vedic chants, mantras, and bhajans as a type.

Prepared August 2026	Audience Site owner / factory brief	Status Implemented on this codebase
--------------------------------	---	---

1. Ideal AI framework — for this factory

Do not introduce LangChain, LlamaIndex, or a Next.js app router for this agent. GenAiMakers already has a working pattern: the browser never sees a key; PHP on Hostinger proxies to Groq (Trolley) or Anthropic (House & Land); deterministic code owns the numbers and the rules; the model writes language. Sloka Agent copies that shell.

What you already run

- **Front end:** static HTML + DM Sans / Instrument Serif, same header and products IA.
- **API:** PHP [api/*.php](#) with CORS allow-list, JSON bodies, file-or-DB rate limits.
- **Models:** Groq chat completions (JSON schema) and Anthropic Messages. Keys live in [api/config.php](#) / env.
- **CSP:** [connect-src 'self'](#) — the browser cannot call Wikipedia or Groq directly. Search must be server-side.

Recommended architecture (what we shipped)

A small **planner in PHP**, not a graph framework. One POST to [/api/sloka.php](#) (legacy [/api/bhajan.php](#) still proxies here):

- **Parse** type (sloka / mantra / Vedic chant / bhajan), deity, language, mood, tradition from free text plus optional form fields.
- **Safety gate** before any tool or model call (copyright-copy, sexual, child, hate, mockery, politics, scams, scripture dumps, restricted diksha mantras). Public traditional mantras are in scope; Gurbani is search-only.
- **Search tools** (server curl): Wikipedia, Internet Archive, DuckDuckGo HTML. Return titles, URLs, source, language, attribution. Truncate snippets. Never fetch or store audio. Never paste lyrics.
- **Rank / filter:** a Wikipedia page about “Sloka” or “Mantra” alone is related reading; an Archive.org item or a page titled as a sloka/mantra/chant/bhajan is an existing listing. Generic YouTube and Archive *search pages* are catalogues, not hits.
- **Generate only if listings are empty** (or the user asks to compose anyway). Groq with a strict JSON schema, then `json_object`, then a smaller model if the large one is unavailable. If the Groq key is missing, the on-site composition factory writes original verses from a multilingual bank (sloka, mantra, Vedic-style chant, bhajan).

Multilingual generation is Groq’s job when the key is present (Telugu, Hindi, Sanskrit, Tamil, Kannada, Malayalam, Bengali, Gujarati, Marathi, Punjabi, Odia, English, Bhojpuri, Awadhi). The factory bank is the fallback so the page still works on a laptop without PHP keys.

What to add later — only if volume justifies it

- Brave Search or Google Programmable Search for cleaner discovery (still links only).
- YouTube Data API for video metadata — still no downloads, no embeds of full tracks if you want to stay conservative.
- Optional spoken demo via a licensed TTS vendor, as a Pro add-on, never cloning a famous singer.
- A Python LangGraph service only if you outgrow PHP and already have a second runtime. It is not required to ship.

Ideal stack in one line: Hostinger PHP planner + Groq JSON generation + Wikipedia/Archive discovery tools + the existing GenAiMakers factory UI. That is what can actually go live this week.

2. Safety measures — implemented, not only promised

Religious generation is a trust product. The agent is useful only if families can open it without fearing a copied film song, a mocked deity, or a scraped MP3.

Copyright and discovery

- Search is **discovery**: metadata and outbound links. Snippets are capped and long quoted verse blocks are stripped.
- Lyric-dump and pirate hosts are skipped in DuckDuckGo results.
- Requests of the form “full lyrics of...”, “download the mp3”, “torrent” are refused. We will find public links or write something new — we will not transcribe.
- Generated output is labelled **generated original**. Listings are labelled **found existing** or **public catalogue**.
- The model is told the titles just found so it must not paraphrase those works.
- Raga names are suggestions for a singer, not a claim to a recording.

Sacred content and respect

- **Tone is bhakti, not comedy.** Mockery, roast, and “make fun of [deity]” are refused.
- No hate, no calls to violence, no sectarian campaign songs.
- **Public traditional mantras are in scope:** Gayatri, Mahamrityunjaya, Om Namah Shivaya, common stotras, public Vedic shanti mantras. Search finds encyclopaedia and catalogue pages (links only).
- **Still blocked:** restricted diksha / initiation-only / guru-secret mantras, and “give me the full initiation mantra of [living guru/sect]”.
- **No scripture dumps:** we will not paste entire Vedas, Upanishads or Gita. Short original slokas are fine; if a public source exists, we link it.
- **Gurbani:** we do not compose Sikh scripture. Linking to public catalogue entries for existing Gurbani is allowed as found existing with attribution. Punjabi Hindu-style output is labelled **not Gurbani**.
- Generated Vedic-style chants are labelled **original compositions in Vedic style**, not authentic Veda recitation from a specific shakha — unless the hit is a found-existing public source.
- Original verses are explicitly **not liturgy** and not a replacement for a temple’s authentic practice.

Content, children, scams, abuse

- Sexual content is refused. Child-sexual content is refused with no alternative.
- Political propaganda and conversion campaign songs are refused.
- Scam patterns (send money, crypto giveaway, lottery) are refused.
- Off-scope crime help is refused. This agent only finds or composes slokas, mantras, Vedic chants and bhajans.

Platform hygiene

- Input sanitisation: strip tags, cap 800 characters, same CORS allow-list as the rest of /api.
- Rate limit bucket **sloka**: 24 searches / hour, 12 forced compositions / hour, token-or-IP keyed like Trolley.
- No model key in the browser. Groq is server-side only.
- CSP remains connect-src self; the page does not open a third-party model from JavaScript.
- Privacy policy now names Sloka Agent: prompts are not stored as a lasting profile; lyrics and audio are not kept.
- On-page fallback if PHP is down: local original composition + public catalogue links, with an honest status line.

Human review still applies if you white-label this for a temple or a radio station. The public demo is self-serve and labelled. Anything that goes out under another organisation’s name should be reviewed before it is printed or broadcast.

3. How this generates revenue

Trolley already has Stripe price IDs in config (annual, monthly, student). Sloka should reuse that billing spine rather than invent a fourth payments vendor. Sacred text should not carry display ads on the generated song itself.

Free layer (proof of factory — what is live now)

- No account. Search public catalogues. A small number of original compositions per hour (rate-limited).
- Purpose: show the factory, collect discovery-call leads, and seed diaspora and temple traffic.

Sloka Pro (same Stripe as Trolley Pro)

- **Credits:** 30 original compositions / month, then pay-as-you-go packs for festival weeks (Janmashtami, Navaratri, Ganesh Chaturthi).
- **Keepers:** save a booklet of family originals (PDF), share a private link, remember preferred language and deity.
- **Voice (later):** licensed TTS in Hindi/Telugu/Tamil — never a celebrity voice clone. Price this as an add-on; it is the expensive SKU.
- **Festival packs:** a week of morning/evening originals for a household or a small mandir, billed once.

B2B factory briefs (the real margin)

- **Temple / ashram white-label:** their domain, their deity calendar, their human review queue, no public compose of rival traditions if they do not want it. This is a signed brief — the same motion as aged care.
- **Community radio and diaspora associations:** a weekly original in one language, with a discovery playlist of public-domain Archive items.
- **Language houses and schools:** classroom-safe originals with transliteration and meaning — already in the artefact.
- Point these at [contact.html?topic=Sloka Agent](https://genaimakers.com/contact.html?topic=Sloka%20Agent) and the 45-minute discovery call.

What not to do

- Do not run banner ads against generated lyrics. It cheapens the offering and will attract the wrong partners.
- Optional, clearly marked sponsorship on the **catalogue** panel only (“public sources”) is acceptable later; keep it off the original poem.
- Do not sell other people’s recordings. Discovery links send traffic to Wikipedia, Archive.org and YouTube; you are not a streaming service.

Practical 90-day mix: keep the public agent free (factory proof). Add Pro credits on the existing Stripe prices. Book two temple/community discovery calls from the Sloka page CTA. That is enough to see whether this is a product or a lead magnet — both are valid for a factory.

4. Where it sits on genaimakers.com

The site's information architecture is deliberate: top nav is Products / Industries / How we build / About. Live tools are not extra nav chrome; they are proof cards. Sloka Agent follows that rule.

Placement we implemented

- **Home** — third live-agent card in “Three agents you can use today”, hero tab, footer Live agents, Culture industry tile. Stats updated from 2 to 3.
- **Products** — third live card, saffron/gold treatment so it is distinct from Trolley's teal and House & Land's gold-land.
- **Industries** — new **Culture** tile (live). This is the honest industry home: not a fake “DevotionAI” suite.
- **How we build** — third “see it on...” button. The factory steps still hold: intake, tools, artefact. Search is the compute; the verse is the artefact.
- **About, Contact, 404, Privacy** — wired. Contact topic “Sloka Agent” for white-label briefs.
- **Not in the top nav.** Adding a fifth primary link would dilute the factory story. People reach it the same way they reach Trolley.
- Short URL [/sloka](#) rewrites to the agent page. [/bhajan](#) and [bhajan-agent.html](#) remain aliases.

Why this is the right slot

House & Land proves high-stakes Australian compute. Trolley proves a worldwide household artefact. Sloka Agent proves multilingual, tool-using generation with a conscience — search first, compose second — and opens a culture/community brief line the factory did not have. It is still one job, one page, live today.

How to run it

- Upload the new HTML/CSS/JS, [api/sloka.php](#), [data/sloka-corpus.json](#), and this PDF. Keep the [/bhajan](#) aliases.
- Ensure [GROQ_API_KEY](#) is set (same as Trolley). Without it, original verses still come from the factory bank for sloka, mantra, Vedic chant and bhajan.
- Open [/sloka-agent.html](#), try “Krishna sloka in Telugu”, “Gayatri mantra”, “Vedic chant for peace”, and “Krishna bhajan in Hindi”. Confirm catalogues appear, then an original if no specific listing exists.
- Optional: `php api/sloka.php --self-test` on the server.

Owner decision, not a blocker: keep Sloka as a free proof agent (recommended while you take temple briefs), or attach it to Trolley Pro credits in the next Stripe pass. The product is already in the IA either way.

GenAiMakers is a factory. This agent is live because it was built — not because a catalogue needed another logo.