Naomi Hosting — An Offering Concept

Managed hosting and mobile build-out where a provably-owned, values-formed, governed AI is the underpinning — not a rented wrapper over someone else's model. This is a real offering concept, ready to build with the right partner. Written 2026-07-04; revised 2026-08-10.

The idea, in one line

Managed web hosting — and mobile app build-out — where a *provably-owned, values-formed, governed AI* **is the underpinning.** Sites and apps built, hosted, and maintained on Cloudflare, with Naomi (owned model), PII protection, and an attribution/audit trail baked in, backed by two granted patents. Not a reseller of OpenAI; the AI is ours. A parish, a diocese, a hospital system, or any values-aligned institution does not need two vendors for this — one relationship covers the website, the member-facing app, and everything AI touches in both.

Why it could matter — the market gap

The market has the *pieces* but not this combination (checked the web, July 2026):

- **"AI-managed hosting"** (Webbee, Hosting.com) — the AI *runs the hosting* (a management copilot), it isn't the underpinning of the client's app; models are vendor-backed.
- **White-label AI for agencies** (HighLevel, Vendasta) — the reseller *doesn't own the model*; a wrapper over OpenAI/Anthropic.
- **Enterprise "sovereign / private AI"** (Safe Swiss Cloud, Cloudera, gateways) — heavy infra plays where the *client brings their own model*; even "sovereign AI" usually means self-hosting an open one.
- **Cloudflare's own AI** — the AI Gateway *proxies* OpenAI/Anthropic; Claude Managed Agents run *inside* Anthropic. Vendor-backed, not owned.

The gap: no one found offering *managed hosting + build/maintain (agency)* where a **provider-owned, values-formed, governed model** is the underpinning, with **PII protection backed by granted patents** and a **native attribution trail** — for SMB-to-mid-market and values-aligned institutions.

The moral floor is the differentiator, not a feature on a list

Every AI vendor in this space is selling capability. None of them are selling **conscience.** A reseller of OpenAI cannot promise a client that the model underneath their site will never help produce what harms a person, because the model isn't theirs to govern — the guardrails belong to a company in San Francisco that can change them without asking. Naomi's moral floor is not a setting. It is written, it is consistent, and it does not move because a vendor updated a policy overnight. The same is true for any institution, of any tradition, that wants an AI it can actually govern — hospitals, schools, financial and legal practices, any values-aligned organization that cannot accept "trust us" from a vendor whose values are unstated. This is not a niche offering for one kind of client. It is the offering for anyone who wants this architecture.

Why *our* AI is what makes it differentiated, technically

A reseller of OpenAI can't make these promises; we can, *because the model is ours*:

- **Client data never leaves to a frontier vendor** — inference on the owned model; prompts/PII don't go to OpenAI/Anthropic.
- **Provable PII protection, not asserted** — the patents *are* the architecture: **USP 11,042,641** (deployment/ inference governance, audit ledger, revoke-on-violation, no-PII-insights) + **USP 11,151,254** (user-defined, purpose-limited, audited data access + cryptographic erasure). A *different* patent class than the market's PII players (whose patents are on *classification*).
- **Attribution/audit native** — every AI action ledgered by reference (data-minimized) — compliance evidence by default.
- **Margin + low-ops** — Cloudflare carries scale/uptime/DDoS/SSL; maintenance ≈ content/feature changes (Naomi Code does those); no per-token vendor cost passed through.

The stack

Cloudflare (edge / compute / static / Workers) + **Supabase** (Postgres / auth / realtime — already used by Enoch & IMDM) + **Naomi** (the owned AI) + **the patents** (the PII/governance proof). Light DB work (schema, migrations, RLS policies) is a clean service line. *(PII caveat: Supabase is a vendor and PII lives there — for sensitive tenants use Postgres RLS + encryption, or D1/owned, plus the 254 posture, so "PII protection" stays true, not a label.)*

The market segments — split by whether the AI is on the hot path

The load risk only bites the AI inference path; everything else scales on Cloudflare without limit. The same split governs mobile: a companion app for scheduling, giving, and messaging is static-to-light; a concierge-style AI chat inside the app is not, today.

- **Static / low-AI sites and apps — buildable today, any scale, zero Naomi load.** Institutional sites and companion apps — parishes, dioceses, schools, nonprofits, hospital systems; small-business and professional sites (law, medical, financial). AI is a selling point and a light feature, not the hot path.
- **Light-AI sites and apps** (chatbot, doc Q&A, search, push-notification intelligence) — owned model at low concurrency, or the Workers AI bridge.
- **Heavy-AI, high-concurrency sites and apps** — custom scope; not offered today.
- **Mobile build-out — its own line, independent of the AI question.** A native or cross-platform app (React Native / Flutter) is a build-and-maintain service like the website, priced on its own scope. AI features inside it follow the same split above.

What the market charges (checked August 2026, sourced)

Real numbers, both ends of the market:

The commodity end — what a vendor-backed wrapper actually costs:

- Church / nonprofit website: custom agency design $3,000–$15,000+ one-time, hosting/maintenance $85–$175/month.
- Small-business website maintenance: $95–$195/month full-service.
- AI chatbot as an add-on: $150–$450/month all-in, with the market's own reporting naming hidden fees as routine — real cost lands near 2.3× the sticker price.

What AI hosting specifically charges — this is the real comparable, not generic enterprise software:

- A defined-scope private AI/LLM deployment (architecture, deployment, integration, stabilization): builds published from **$8,000**, before any ongoing hosting.
- Production-scale AI hosting — a private model actually serving real traffic, not a demo — runs from **$10,000/month** and climbs past **$100,000/month** for large multi-GPU configurations.

- Institutional website builds at this tier (100+ pages, compliance, integrations) run **$100,000–$250,000+**, with maintenance at 15–20% of build cost annually.
- Agency-built mobile apps average **$90,000**.

This offering charges into that range, not toward the commodity floor. None of the commodity-end numbers buy what this is built on: every one of them is a vendor-backed model, a wrapper over OpenAI or Anthropic, or a platform with no attribution trail and no patent behind its privacy claim. **The patents and the moral floor are the reason to price as premium, proven governance, not as a website with an AI feature bolted on.** This is also consistent with 33 Consulting's own established rate class for differentiated work — $15K–$150K per engagement — not a discount arm of the firm.

What we charge

A single price per tier, not a range — this is what it costs.

Tier	What it is	One-time build	Monthly
Foundation — governed site	An institution-grade site on the owned stack, patent-backed governance and attribution from day one	$45,000	$1,500/mo
Governed AI	The above, plus Naomi (the owned model) live on the site — chat, search, guided content	+$40,000	+$7,500/mo
Heavy-AI	Custom scope, high-concurrency — not offered today	—	—
Mobile build-out	A companion app on the same governed stack, cross-platform	$95,000	$1,500/mo

The financial model — what it takes to run this

The infrastructure cost is close to zero. Cloudflare carries scale, uptime, DDoS, and SSL for every client on the same platform at no added marginal cost. The model is already built and owned — no per-token vendor bill on any client's traffic. The real cost of this business is not compute. It is people.

Revenue per client, at the prices above:

- A Foundation-only client: $45,000 build, $18,000/year recurring ($1,500/mo).
- A Foundation + Governed AI client: $85,000 build, $108,000/year recurring ($9,000/mo).
- A Mobile build-out, add-on to either: +$95,000 build, +$18,000/year recurring.

What each path needs, people-wise:

- **Foundation only, 1–3 clients:** no dedicated hire. The build is a project engagement; ongoing support is covered directly, inside the relationship, the way any consulting engagement is.
- **Foundation only, 4+ clients:** $72,000+/year in recurring revenue alone — enough to fund one dedicated account/support role. A technical account manager runs $90,000–$120,000/year fully loaded (2026 US market data); four Foundation clients' recurring revenue alone covers most of that, and the build fees cover the rest in year one.
- **Any Governed AI client:** requires a second, technical skill set — someone who can own the inference tier: load, uptime, and the model itself, not just content changes. One Governed AI client's recurring revenue ($108,000/year) funds this role on its own.
- **Mobile build-out clients:** the build fee funds the engagement directly; ongoing maintenance revenue ($18,000/year per client) is additive to whichever support role above is already in place — it does not, by itself, require a third hire until volume justifies it.

The architecture is already right. What this model shows is that the business does not need volume to work — it needs the first few clients at this pricing, and the second hire funds itself from client four onward. That is a real plan, not a hope.

This is what it looks like to build AI on an owned foundation — not rented, not asserted, provable.