



ChatGPT Ads Agency

Request for Proposal (RFP) Template

A structured way to evaluate any ChatGPT Ads agency — before you sign anything. Free to use, whether or not you choose to work with us.

How to Use This Template

Fill in Section 1 with your own business details. Send Sections 2 through 8 to every agency you're evaluating — the same questions, to each, so answers are actually comparable. Use Section 9 yourself, after responses come in, to score each agency side by side instead of relying on general impressions.

This template works regardless of which agency you end up choosing. It exists because most ChatGPT Ads pitches right now are long on confidence and short on specifics — this is built to surface the difference.

Section 1 — Your Business Overview

Fill this in once, then attach it to every RFP you send.

Business name: _____

Industry / category: _____

Primary location(s) served: _____

Current AI visibility status (if known): _____

Products or services to be advertised: _____

Target timeline for readiness work: _____

Section 2 — Scope of Work

Q1. What specific readiness deliverables will you provide before ChatGPT Ads launches in India?

Why this matters: A real answer names deliverables — an entity readiness report, confirmed account setup, creative assets, a product feed if relevant. A vague answer talks about "strategy" and "campaigns" without specifics, for a channel that isn't running yet.

Q2. What exactly is included in your fee, and what counts as a separate cost later?

Why this matters: Readiness work should be priced and scoped separately from future campaign management. If a quote bundles the two together, ask why — there's no live campaign to manage yet.

Section 3 — Technical Capability: Account & Entity

Q3. Who creates the OpenAI advertiser account — us or you?

Why this matters: Correct answer: you do. OpenAI's signup flow requires the client to create the advertiser account first; the agency is added afterward as a team member. If an agency says they'll create or host the account for you, that's a structural misunderstanding worth flagging immediately.

Q4. How will you audit our current entity and brand data across sources?

Why this matters: Listen for specifics — schema markup, cross-source consistency checks (is your business name, address, and category the same everywhere it appears), structured data. "We'll do some SEO" is not an answer to this question.

Q5. Can you show a named case study of entity or schema work you've completed?

Why this matters: Anonymous case studies ("a growing consumer brand," "an enterprise client") are unverifiable by design. In a field this new, agencies with genuine experience should have a small enough client list that naming results isn't a problem.

Section 4 — Platform Mechanics Understanding

Q6. Which of our product or service categories are eligible for ChatGPT Ads, and which face restrictions?

Why this matters: Healthcare and medicine, legal services, most financial services, gambling, and political content are excluded outright, regardless of market. An agency that hasn't checked this against your specific business hasn't done the homework.

Q7. How does targeting work on this platform, and how is it different from keyword-based advertising?

Why this matters: Correct answer references conversational context, not keyword matching. If the explanation sounds identical to how they'd describe Google Ads, that's a red flag.

Q8. What ad formats will our creative need to support, and how are you preparing for them?

Why this matters: Formats include standard text-based placements and, as the platform matures, product carousel sequences for retail/D2C. An agency should be able to name the difference, not describe one generic "ad creative" process.

Section 5 — Attribution & Tracking

Q9. How will you track conversions when there's no clean "last click" to attribute to?

Why this matters: Conversational ad journeys don't follow the linear click-path that standard last-click attribution assumes — a user can see a sponsored card, keep chatting across several turns, and convert without a single trackable click. A capable agency should reference server-to-server or Conversions API (CAPI) based tracking. "We'll add a pixel" is a Google Ads answer to a ChatGPT Ads question.

Section 6 — Deliverables & Timeline

Q10. What are the specific, dated milestones for readiness work?

Why this matters: Not a campaign launch date — India's timeline isn't confirmed by OpenAI. Look for phased milestones instead: entity readiness by one date, account and creative readiness by another.

Q11. What does "done" look like for the readiness phase, separate from future campaign performance?

Why this matters: Since no campaign exists yet, "success" for this phase should be readiness completion — entity score improvement, account verification, creative sign-off — not projected ROI on a channel that isn't live.

Section 7 — Budget & Pricing Structure

Q12. Is your fee a scoped project fee for readiness work, or an ongoing retainer? Please justify.

Why this matters: A retainer priced as if a campaign were already running is a fee for a service that doesn't exist yet in India.

Q13. If you're quoting a specific campaign budget or CPC/CPM figure for India, what is it based on?

Why this matters: Nobody has confirmed India pricing yet. A credible agency references live-market data (currently the US, Canada, Australia, and New Zealand) as a directional ballpark and says so plainly — not a number presented as confirmed.

Section 8 — Case Studies & References

Q14. Please provide 2–3 named client references we can contact directly.

Q15. What AI-visibility or GEO results have you achieved for a client in our industry, or a comparable one?

Section 9 — Evaluation Scorecard

Use this after responses come in. Score each agency 0–5 per row based on their written answers, then compare totals. This is for your internal use — not something you send to agencies.

Evaluation Criteria	Agency A (/5)	Agency B (/5)	Agency C (/5)
Understands account ownership (client creates, not agency)			
Entity/schema audit capability — specific, not generic "SEO"			
Named, checkable case studies provided			
Category eligibility knowledge (excluded categories correctly identified)			

Evaluation Criteria	Agency A (/5)	Agency B (/5)	Agency C (/5)
Attribution/tracking understanding (CAPI, not just "a pixel")			
Pricing transparency — readiness fee scoped separately from campaign budget			
Realistic timeline (phased milestones, not a fixed launch date)			
Overall communication clarity			
TOTAL (out of 40)			

A
strong
agency

should score well across every row, not just the ones that flatter their pitch. An agency that scores high on communication but low on the technical rows (3, 5, 6) is likely better at marketing itself than at the actual work.

Provided by KickAss Digital Marketing — India's AI-First ChatGPT Ads Agency

kickassdigitalmarketing.com · Goa HQ, serving businesses across India