

The Agent-Readiness Checklist

12 signals that make your business visible — and callable — by AI agents.

Why this matters now

AI agents are replacing traditional search. ChatGPT, Perplexity, Gemini, and Google AI Overviews answer questions directly — 68% of Google searches now end without a click (SparkToro x Similarweb, 2026). Where an AI Overview appears, organic click-through has dropped 61% (Seer Interactive, Jun 2024–Sep 2025: 1.76% to 0.61%), while brands that ARE cited in the AI Overview receive 35% more organic clicks than those that aren't.

In an Ethereum Forge scan of 100 local business sites, 93% were not agent-ready: no structured data, no machine-readable summaries, no way for an AI agent to understand — let alone cite or call — the business.

Sources: Seer Interactive, AIO Impact on Google CTR, Sep 2025 update · SparkToro x Similarweb zero-click study, Jan–Apr 2026 (US) · Ethereum Forge scan data (first-party).

The 12-point checklist

1. Schema.org JSON-LD markup

Organization, Service, and FAQPage structured data in JSON-LD. This is the primary way AI agents understand what your business is, what it offers, and how to cite it.

2. llms.txt and llms-full.txt

Plain-text summaries of your business at /llms.txt, written for language models. Emerging convention already fetched by AI crawlers.

3. An agent card at /.well-known/agent-card.json

A machine-readable card (A2A convention) declaring who you are and what structured actions, if any, agents can take with you.

4. robots.txt that addresses AI crawlers

Explicit User-agent rules for GPTBot, ChatGPT-User, ClaudeBot, PerplexityBot, and Google-Extended — so AI engines know they are welcome.

5. XML sitemap

A current /sitemap.xml so crawlers and agents can enumerate every page you want discovered.

6. Clean titles, meta descriptions, and canonical tags

One canonical URL per page, a 50-60 character title, and a 120-160 character description. AI answers quote these surfaces directly.

7. Semantic HTML structure

Real , , , , and landmarks. Agents parse structure, not pixels.

8. Server-rendered content (not SPA-only)

Your content must exist in the HTML response itself. A JavaScript-only shell is invisible to most AI crawlers.

9. Structured contact paths

mailto: and tel: links agents can extract reliably — not contact info trapped inside images or scripts.

10. Social/share images and alt text

Open Graph and Twitter card images plus descriptive alt text on at least 80% of images — multimodal engines read both.

11. Breadcrumb and speakable schema

BreadcrumbList markup for site structure and SpeakableSpecification for voice assistants — low-effort, directly machine-read.

12. WebMCP declarative forms (draft standard)

WebMCP is a draft web standard (W3C Web Machine Learning Community Group) that lets a site declare callable tools to browser agents. Deploy it alongside the stable signals above — not instead of them.

Check your own score in 60 seconds

The free Agent-Readiness Scan checks nine of these signals automatically and returns a grade, a score, and your top findings — no account, no credit card, no sales call. Run it at etherealmedia.ai/scan.