AI Citation Service: Generative Engine Optimization Architecture Brief

Tygart Media Technical Reference · Grounded in Origin-Log Forensics & Multi-Engine Citation Testing (Nov 2026)

> **What the Service Delivers** — The **AI Citation Service** restructures commercial and technical web pages into standalone answer packets, upper-40% structured comparison matrices, dual-path Markdown twins (`<link rel="alternate" type="text/markdown">`), 2-page House Reference PDFs, and verified Schema.org graphs so ChatGPT, Perplexity, Claude, Microsoft Copilot, and Google AI quote your business.

1. Why Most Ranked Pages Still Lose AI Citations (The RNC Problem)

Ranking in traditional search indexes is necessary but no longer sufficient for AI visibility. Across live multi-engine audits, pages frequently rank #1–#2 in Bing or web search retrieval yet receive **zero inline footnote citations** (**Retrieved, Not Cited / RNC**) because narrative prose buries the direct answer below long hero sections or client-side widgets.

Optimization Dimension	Traditional SEO Focus	Tygart Media AI Citation Engineering	Measured Structural Impact
Opening Paragraph (`p[0]`)	Narrative hook / keyword repetition	**45–55 word** standalone declarative answer packet	Matches LLM extraction window; eliminates `<30w` snippet stubs
Structured HTML Data	Single late-page summary or image	**>= 2** upper-40% HTML tables (`table-layout:auto`)	Directly feeds multi-cell comparison parsing without `has-fixed-layout` clipping
Machine-Readable Twins	HTML only; unverified `/llms.txt`	LF `.md` twin + `<head>` `<link rel="alternate">` tag	Verified direct Perplexity & Copilot agent citations on `.md` and `.pdf` formats
Downloadable Spec Artifacts	None or gated lead-trap PDFs	Un-gated **2-page House Reference PDF** (`useA85=0`)	Indexed independently across Bing/OpenAI PDF retrieval lanes
Indexation & Crawl Hygiene	Manual spot-checks	Headless GSC Sitemaps + `IndexNow` + `NoRedirect HTTP 200`	Removes `301` redirect latency tax (`up to 37.56%` on live fetchers)

2. Five-Layer Citable Page Blueprint

- **Layer 1 — Lead Answer Packet (`p[0]`):** One concrete `45–55` word summary stating pricing tiers, edition scopes, limits, and primary source provenance.
- **Layer 2 — Upper-40% Comparison Matrices:** Clean `<thead>` + `<tbody>` HTML tables placed above the 40% scroll mark with consistent units and primary-doc citations.
- **Layer 3 — Synced Markdown Twin (`/slug.md`):** UTF-8 LF twin advertised via `<link rel="alternate" type="text/markdown">` in `<head>` (zero JS required).
- **Layer 4 — Companion 2-Page Reference PDF:** Text-selectable vector PDF with clear H2 sections and source attribution.
- **Layer 5 — Schema & Crawl Proof:** Valid `application/ld+json` (`Service`, `Article`, `FAQPage`), `<meta name="description">` `142–153` chars, and `NoRedirect HTTP 200` canonicals.

AI Crawler Architecture & Service Verification Scorecard

Page 2 · Primary Bot Documentation Mapping & Client Deliverable Contract

3. Search Indexers vs. Live Prompt-Time AI Fetchers (Primary Docs)

Platform / Engine	Search Indexer UA	Live Fetcher UA	Primary Crawler Rule & Technical Requirement
OpenAI / ChatGPT Search	`OAI-SearchBot`	`ChatGPT-User`	`developers.openai.com/api/docs/bots`: Allow `OAI-SearchBot` in `robots.txt`; live `ChatGPT-User` reads raw HTML (`~24h` update lag).
Anthropic / Claude	`ClaudeBot` / `Claude-SearchBot`	`Claude-User`	`support.claude.com/en/articles/8896518`: Independent `robots.txt` controls for training (`ClaudeBot`) vs. user-triggered retrieval (`Claude-User`).
Perplexity	`PerplexityBot`	`Perplexity-User`	`docs.perplexity.ai/docs/resources/perplexity-crawlers`: Surfaces search citations cleanly from static HTML, `.md` twins, and PDFs.
Microsoft Copilot / Bing	`bingbot`	Copilot substrate	Stricter `RETRIEVED_NOT_CITED` gate on single-table pages; favors `>= 2` upper structured tables, `.md`, and `IndexNow` freshness.
Google AI Overviews / Mode	`Googlebot`	`Google-Extended`	Strict **GSC indexation gate**: pages stuck in `Crawled/Discovered - currently not indexed` score `0` regardless of on-page quality.

4. Engagement Workflow & Audit Receipts

1. **Baseline AI Citation Check (Free):** Multi-engine audit showing exactly where ChatGPT, Microsoft Copilot, Perplexity, and Google AI currently mention or skip your business.
2. **Rival Autopsy & Structural Spec:** Side-by-side comparison of the domains currently winning citations on your target commercial queries (`WebGuide` vs. official docs vs. PDFs).
3. **Production Execution (**`HTML` **+** `PDF` **+** `.md` **+ Schema):** Full deployment of validated answer packets, HTML matrices, PDF briefs, Markdown twins, and Schema.org JSON-LD.
4. **Verification Receipts:** Live `NoRedirect HTTP 200`, Playwright desktop (`1440w`) + mobile (`375w`) layout checks, and monthly Bing/GSC citation & performance reporting.

> **Next Step** — Request a free AI-citation check at `will@tygartmedia.com` or review the live service specification at `https://tygartmedia.com/ai-citation-service/` ($497 standalone 5-day AI Visibility Audit available separately; project rebuilds scoped post-audit).