TYGART MEDIA

AI SEO Vendor Test Quick Card • The Five Receipts Every Service Should Produce

Current version — Nov 1, 2026. Update history: https://tygartmedia.com/top-rated-ai-seo-services-small-businesses-test/

THE GOVERNING RULE: RATE THEM ON EVIDENCE, NOT ADJECTIVES

You searched for top rated AI SEO services for small businesses and found listicles written by the firms on the list. We are one of the firms you could hire, so this card is not a neutral ranking. It is something far more useful: **the five receipts every service on your shortlist must produce**, the **three red flags that end a sales call**, and the worked benchmarks to verify them. Bring this test to every pitch you sit through. Including ours.

Receipt 1: Their Own Citation Trend

TREND VS SNAPSHOT

The Call Question: "Is an AI engine citing *your* site more this month than last?"

Standard Required	Proof Instrument	Vendor Test & Failure Mode
Drowning in their own visibility: A firm selling visibility must have proof of its own active citation base.	**Engine's own reporting:** Bing Webmaster AI Performance reports (free, auditable). Never accept a proprietary dashboard.	**Trend, not snapshot:** A total climbing across monthly windows is a practice; a single quoted total is a photograph of weather.

Receipt 2: Grounding Queries in Your Vertical

VERTICAL GROUNDING

The Call Question: "Which questions in *our* industry does the engine already ground on your clients' pages — and on your own?"

Standard Required	Proof Instrument	Vendor Test & Failure Mode
Footprint where buyers live: Citation counts on random topics prove nothing about your market.	**Dated Query Export:** Bing AI Performance query-level raw export listing exact industry grounding queries.	**Make them show it:** A firm cited for contractor certification or insurance claims has vertical authority. Wins in software pricing mean footprint is somewhere else.

Receipt 3: One Climbed Page, With the Export

BEFORE / AFTER PAPERWORK

The Call Question: "Name one page you raised, and produce the two exports: before and after."

Standard Required	Proof Instrument	Vendor Test & Failure Mode
A real climb survives inspection: Every service has a story about a page they raised. Demand the receipts.	**Dated, downloadable exports:** Must show exact URL, raw citation total, and specific trailing reporting window.	**Stories without paperwork = marketing:** If they cannot produce before/after exports, you are about to pay retainer rates for marketing.

Worked Example Benchmark (Tygart Media Report Card): October 15, 2026 Bing export: **1,459,726 citations** across **922 pages** (up from 269,399 across 324 pages on Sept 27). Grounding shares: 24.50% "claude pricing", 33.85% "claude ai pricing", 54.46% "anthropic console". Client vertical climb: 1,635 → 1,913 → 2,454 → 2,790 citations while duplicate consolidated from 304 to 6.

Receipt 4: How They Measure

METHOD VS METRICS

The Call Question: "How will you report *our* results? Walk us through the exact reporting method, cadence, and data validation."

The Honest Method	Single-Pull Anomaly Trap	Audit Threshold
Dull, consistent cadence: The same reports, the same trailing window, read against each other once a month.	**Citation share is a rate, not a switch:** A single data pull can contain cells that are simply wrong (e.g. 2,796 citations at 0.00% share).	**Caught tools lying:** A second report must settle anomalies. A firm that has never caught its tools lying has not been measuring long.

Receipt 5: They Publish Their Own Numbers

PUBLIC REPORT CARD

The Call Question: "Where is your public report card on your own website that we can inspect and re-check next quarter?"

Signal Cost	Engine Shortlists	Evidence vs Adjectives
Costs the firm something: A public record carries reputational risk if numbers stall.	**Real rating is assembled by engines:** AI engines build buyer shortlists from actual footprints, not landing pages.	**Firms that publish get rated on evidence:** Firms that don't publish are asking you to rate them on adjectives.

THREE RED FLAGS THAT END THE SALES CALL IMMEDIATELY

Red Flag 1: A visibility score with no source report behind it

If the number cannot be traced back to an engine export you can pull yourself, it is a mood, not a measurement. Proprietary black-box scores hide lack of real engine citations.

Red Flag 2: A promise to "get you cited in ChatGPT"

No named instrument exists for how either of you will know. ChatGPT and Claude synthesize answers from footprints that cannot be totaled or tracked via verified public reporting dashboards.

Red Flag 3: Case studies with percentages but no pages

Growth claims that never name the URL, the query, or the trailing window are unfalsifiable by design. Each of these flags tells you in advance exactly how the firm will report your results later.

Report Card Required Limits: Why Limits Prove Authenticity

A report card without limits is an ad. Every credible vendor submission must state: **1. One engine family:** Bing/Copilot ledgers are auditable; ChatGPT/Claude cannot be totaled the same way. **2. One showcase, one niche:** Distinguish agency tech/software coverage from client industry verticals. **3. Trailing measure, not a promise:** Citation shares shift as engines re-ground; last quarter's climb guarantees nothing about next quarter. **4. AI Governance:** How the firm keeps AI agents from publishing unchecked work is part of what you are actually buying.