

FREE WORKSHEET · INTEGRATED.SOCIAL

AEO Evidence Scorecard

A practical two-page checklist for evaluating transparent AI-search readiness—without pretending that any score guarantees a citation, ranking, lead, or conversion.

How to use this scorecard. For each item, circle one rating: **0 = missing or contradicted**, **1 = partly present or unverified**, **2 = documented, visible, and validated**. Record the evidence URL, audit note, or QA test beside each item. A low score is not a forecast—it is a prioritized worklist.

0–5 · Critical

Resolve eligibility or evidence gaps before making visibility claims.

6–11 · Foundation

Core work exists, but proof, consistency, or measurement is incomplete.

12–16 · Evidence-led

Maintain the evidence register and validate every meaningful release.

Part 1 · Discovery and evidence

Assessment item	Rating and evidence
1. Scope and buyer question clarity Have you documented the priority URLs, target country/language, buyer task, conversion objective, and exclusions?	0 ☐ 1 ☐ 2 ☐ Evidence: [write here]
2. Technical eligibility Do the rendered page, canonical URL, sitemap, internal links, and indexability signals agree? Are important pages accessible without broken routes or blocked resources?	0 ☐ 1 ☐ 2 ☐ Evidence: [write here]
3. Entity and service clarity Can a reader—and a crawler—identify who you are, who you help, the service boundary, proof, author, and next action from the page itself?	0 ☐ 1 ☐ 2 ☐ Evidence: [write here]
4. Evidence and source discipline Are material facts sourced, dated, qualified, and connected to visible content? Have unsupported figures, reviews, prices, and credentials been removed or verified?	0 ☐ 1 ☐ 2 ☐ Evidence: [write here]

Note for regulated or complex topics: add the appropriate legal, clinical, product, privacy, security, or compliance reviewer to the release record. Transparent assessment means making the approval boundary visible—not claiming that a generic checklist replaces specialist review.

Part 2 · Content, measurement, and action

Assessment item	Rating and evidence
5. Answer-first content Does the page answer the core buyer question early, then explain criteria, trade-offs, evidence, limitations, and a practical next step?	0 ☐ 1 ☐ 2 ☐ Evidence: [write here]
6. Semantic and schema parity Are H1–H4 headings semantic and logical? Does the structured data accurately describe visible on-page information, with no duplicate or invented claims?	0 ☐ 1 ☐ 2 ☐ Evidence: [write here]
7. Reproducible measurement Do you separate organic, local, AI-answer, and citation observations? Are prompt, engine, locale, device/session, date, raw output, and cited URLs recorded?	0 ☐ 1 ☐ 2 ☐ Evidence: [write here]
8. Implementation and conversion loop Does each recommendation have an owner, dependency, QA check, release record, and a relevant conversion path for a qualified buyer?	0 ☐ 1 ☐ 2 ☐ Evidence: [write here]

Your 90-day evidence-led action plan

Window	Priority action	Owner and validation
Days 1–15	Fix hard eligibility gates, define the buyer-question set, and assemble the evidence register.	Owner: [write here] QA: [write here]
Days 16–45	Upgrade high-value content, entities, internal links, and conversion routes; publish only source-reviewed claims.	Owner: [write here] QA: [write here]
Days 46–90	Compare pre/post signals in Search Console and the controlled observation log; decide what to continue, revise, or stop.	Owner: [write here] QA: [write here]

Need a defensible plan, not just a score?

Book a no-obligation discovery call to scope your evidence, implementation, and validation plan:

calendly.com/modielnadi