AEO/GEO Visibility Audit Framework

A repeatable methodology for auditing and improving brand visibility across AI answer engines

Néstor G. Trillo | SEO / AEO / GEO Strategist

AEO/GEO Visibility Audit Framework

A repeatable methodology for auditing and improving brand visibility across AI answer engines

Néstor G. Trillo | SEO / AEO / GEO Strategist

> *"First, place the client on the AI-maturity curve. That determines what's even actionable. Then run a fixed question set against the AI engines they actually care about. Score for presence and prominence separately… they move at different rates and mean different things. Turn the gap into a roadmap sequenced to what the team can actually operationalize, not the ideal end state."*

Six Phases

1	**Discover**	Map customer intent first (what the market actually wants), then the client's offerings, prioritized by what matters to them. The client's AI Maturity is scored last: can they serve those offerings to that customer intent at scale?
2	**Audit**	Technical foundation checklist plus a fixed question set. Score for presence and prominence, and replicate sequentially across independent AI-visibility tools, not just one.
3	**Diagnose**	Separate presence from prominence. Rank findings by how much fixing them would move the business, weighed against how hard the fix actually is, not by raw citation count. A "leader" label can still mean weak absolute performance.
4	**Plan**	Sequence tactics to what the client's AI maturity stage can actually operationalize, each tied to a real business outcome. Not a generic content backlog.
5	**Implement**	Ship the highest-priority fix. Validate externally: show the check, don't just claim the change is live.
6	**Measure**	Re-run the fixed baseline on a defined cadence. Separate drift (things changed over time) from impact (our fix caused the change). Conflating the two is a common, avoidable error.

What Makes This Different From a Generic SEO Audit

Presence ≠ prominence. Being cited isn't being the answer. Scored as two separate metrics, every time, not collapsed into one visibility number.

Cross-tool triangulation. A finding confirmed independently by two or three AI-visibility tools carries real evidentiary weight. A single dashboard number, on its own, doesn't.

Denominator discipline. A dashboard percentage is only as trustworthy as what it's actually counted out of. In Case Study 2, one tool's export listed 147 answers, but its real percentages were calculated from only 75 of them. Confirm the actual count against the raw export. Never assume it from the headline number.

AI-Maturity-sequenced roadmaps. Recommendations sequenced to what the team can operationalize now, not the ideal end state a Stage-4 org could execute.

Tool Stack

Screaming Frog / Sitebulb (technical crawl & schema validation) · Profound, Peec AI, and Otterly (AI-visibility & citation scoring, hands-on across all three) · Semrush / Ahrefs / Moz (technical & backlink baseline) · Surfer SEO and AirOps (content structuring & AI content workflows) · GSC / GA4 (proxy signals for lower-AI-maturity clients)

Proven Repeatable, Not One-Off

Run twice, independently, against the same brand (Bank of Hawaii, independent research). Case Study 1: four engines, ten questions. Case Study 2: three engines. Three-tool replication across Profound, Peec AI, and Otterly. A full six-phase run, including an executed technical sample deliverable and a written measurement plan. The same core pattern replicated both times: high presence, low prominence. The brand shows up often, but rarely as the clear answer.