

From SEO Reporting to AI Visibility

A 30-Day Client Brief for U.S. B2B Content Agencies

PUBLIC SAMPLE • PREPARED AUGUST 18, 2026 • RETRIEVED AUGUST 18, 2026 • AI-ASSISTED, HUMAN RESPONSIBLE

Decision owner. Agency Founder / Owner

Decision. What should an agency add to a client’s next 30-day content program without inventing a separate “AI SEO” product?

Executive decision

Recommendation. Add a four-part AI-visibility workstream to the existing SEO/content engagement: (1) crawl and snippet controls, (2) answer-ready evidence and page structure, (3) Google and Bing AI-visibility baselines, and (4) an experiment log that connects visibility to qualified downstream actions. Do not sell guaranteed AI citations or replace foundational SEO.

This is a practical extension of current work, not a new software platform. Google says its AI search features use the same foundational Search requirements and that a page must be indexed and snippet-eligible; Bing’s new AI Performance report measures citations and cited pages across Microsoft AI experiences. OpenAI provides separate crawler controls, so the client’s robots policy should be an explicit business decision rather than an accidental default.

What changed — verified facts

Fact	Verified statement	Agency implication
Google AI eligibility	Google says AI Overviews and AI Mode use Search systems; eligibility requires indexability and snippet eligibility, with no additional technical requirement.	Keep technical SEO, internal linking, people-first content, visible text and structured-data consistency as the base layer.
Google measurement	Google announced dedicated Search Console generative-AI performance reports in June 2026, including impressions, pages, countries, devices and time views.	Establish a pre-change baseline and report AI visibility separately from ordinary click reporting when the client account exposes the feature.
Bing measurement	Bing launched AI Performance in Webmaster Tools public preview in February 2026, including total citations, average cited pages, grounding-query samples and URL-level citation activity.	Add citation and grounding-query review to the monthly diagnostic; treat the data as visibility evidence, not ranking or authority.
Content design	Bing says clear headings, tables, FAQs, evidence and current information can make content easier to reference accurately; it also says no tool guarantees inclusion.	Improve clarity and evidence structure, but prohibit citation guarantees in proposals and reports.
Crawler policy	OpenAI distinguishes OAI-SearchBot (search discovery) from GPTBot (potential model training controls), and crawler access can be managed through robots.txt.	Get client approval for the desired search-discovery versus training posture; document the change and test the resulting crawl behavior.

Inferences — separated from facts

- **Reporting will fragment before it consolidates.** Google and Bing now expose different AI-visibility measures, so a single “AI traffic” KPI would hide material differences. This is an inference from the current product surfaces, not a provider promise.

- **Evidence packaging is becoming part of content operations.** Both providers emphasize crawlability, clarity, freshness and support for claims. A reusable evidence block can reduce later fact-checking and update work.
- **The first commercial opportunity is operational, not magical.** Agencies can package a controlled audit, measurement baseline and structured-content experiment without claiming a proprietary ranking method.

30-day action plan

Week	Action	Output	Evidence
Week 1	Inventory priority URLs; record Google/Bing visibility; review robots, snippet controls, indexability and allowed OpenAI crawlers.	Baseline and control register.	Screenshots/exports, robots snapshot, dated URL list.
Week 2	Select 5–10 high-value pages; add explicit claims, primary-source links, ownership/date, concise answer blocks and consistent structured data where appropriate.	Before/after content pack.	Editorial QA and schema-to-visible-text check.
Week 3	Publish/update; submit discovery signals supported by the client stack; log change dates and target questions.	Change log and experiment register.	HTTP/indexability checks, submitted URL evidence.
Week 4	Compare visibility, citations, qualified actions and client-review time; decide keep, revise or stop for each pattern.	Decision brief.	Metric export, accepted-quality review, next-month decision.

Minimum measurement model

1. **Eligibility:** indexable, snippet-eligible, crawl controls match the client’s policy.
2. **Visibility:** Google generative-AI impressions/pages where available; Bing citations, cited pages and grounding-query samples.
3. **Business action:** qualified form fills, demo requests, newsletter signups or another client-approved outcome after AI/search discovery.

Guardrail. Do not treat citation count as ranking, authority or revenue. Bing explicitly states that its aggregate citation metrics do not indicate a page’s role or authority in an individual answer.

Quality and risk controls

- Keep facts, inferences and recommendations separated.
- Preserve dated source links for all material claims.
- Do not block or allow AI crawlers without a client-approved content-use policy.
- Do not create schema that contradicts visible page content.
- Do not publish unsupported “best,” “leading,” performance or market-share claims.
- Do not report AI visibility as business impact until a client-approved downstream event is connected.

Unverified / Needs Confirmation

Item	Current source	Missing evidence	Judgment	Follow-up
Google report availability	Google launch announcement	Client-account access and property-level availability	Unknown until login	Check during Week 1 baseline; if absent, record combined Web

				reporting and wait for availability.
Bing preview availability	Bing announcement	Client-account access and preview availability	Unknown until login	Check during Week 1; do not infer zero citations from an absent report.
Current OpenAI crawler reach for the client's priority URLs	OpenAI crawler documentation	Server logs, robots result and WAF/CDN behavior	Likely testable, not yet tested	Review robots/WAF and verify crawler access after client approval.

Commercial packaging

- **Promise:** a controlled 30-day baseline, content experiment and decision brief.
- **Do not promise:** citation placement, traffic lift, leads, revenue or a universal “GEO” score.
- **Scope guard:** 5–10 priority pages, one approved content-use policy, one reporting view, one decision session.
- **Renewal trigger:** continue only when the experiment improves a client-approved visibility or downstream signal without increasing senior review time beyond the agreed limit.

Sources

1. [Google — AI features and your website](#) — retrieved August 18, 2026
2. [Google — Succeeding in AI Search \(May 21, 2025\)](#) — retrieved August 18, 2026
3. [Google — Search Generative AI performance reports \(June 3, 2026\)](#) — retrieved August 18, 2026
4. [Microsoft Bing — AI Performance public preview \(February 10, 2026\)](#) — retrieved August 18, 2026
5. [OpenAI — Overview of OpenAI Crawlers](#) — retrieved August 18, 2026

15-minute efficiency-review CTA

If a client-ready brief like this currently requires multiple rounds of source discovery, senior synthesis and fact correction, map one real workflow for 15 minutes. The review will identify whether a four-week, four-brief managed pilot has a credible path to reducing client human time; the 30% target will be measured against an agreed baseline and is not guaranteed.

AI-assisted, human responsible.