

AEO Audit Parameters

Technical Reference Guide

Kabloom Agency Internal Library · Answer Engine Optimisation · Audit Methodology v2.0

6
Dimensions
evaluated

18
Checks
per page

4
AI engines
covered

Table of Contents

	What is AEO and why does it matter?	p. 3
	How AEO ranking works	p. 3
D1	Structured Data	p. 4
D2	Content Clarity	p. 6
D3	Entity Definition	p. 8
D4	FAQ & Direct Answers	p. 10
D5	Technical Signals	p. 12
D6	E-E-A-T	p. 13
	How scores are calculated	p. 15
	Understanding possible inconsistencies	p. 15
	Priority order for fixes	p. 17

What is AEO and why does it matter?

Answer Engine Optimisation (AEO) is the practice of structuring web content so that AI-powered answer engines — Google AI Overviews, Perplexity, ChatGPT Search, Bing Copilot — can confidently extract, summarise, and cite your pages in direct answers.

Unlike traditional SEO, which targets ranked links, AEO targets zero-click citations: the moment an AI engine quotes or paraphrases your page as the authoritative answer to a user's question. For Kabloom clients in the association and congress world, this matters because potential delegates and members increasingly ask AI assistants questions like "What is the best event marketing agency for medical congresses?" — and the answer they receive depends significantly on how well your pages are structured for machine reading — it is one of the most impactful factors you can directly control (AEO structure matters a lot, but domain authority, how established the brand is online, and external links pointing to the page also influence AI citations significantly).

The audit tool evaluates six dimensions, each worth up to 100 points. The **Overall AEO score** is the average of the six dimension scores for the page being evaluated.

How AEO ranking works

AEO is entirely page by page, not site-wide. When an AI engine answers a query, it finds the single most relevant, best-structured page across the web — not the best site overall. So a campaign landing page, a service page, or a blog post each compete independently in AI engine results.

- **1. Audit the pages that matter for each campaign.** A congress registration landing page, a service page targeting a specific association type, or a blog post answering a common delegate question — each competes independently in AI engine results.
- **2. The homepage audit is a brand entity check, not a performance check.** It tells you how well AI engines can identify and describe Kabloom as an organisation, not whether your content ranks for specific queries.
- **3. Campaign landing pages should be audited before launch, not after.** AEO signals are much easier to build in during content creation than to retrofit.



Structured Data

Why it matters

Schema markup is the primary language AI engines use to understand what a page is, who created it, and what it covers without relying on reading and interpreting natural language. Pages without schema markup force AI engines to make educated guesses — and they often get it wrong, misattributing content or failing to surface it for relevant queries.

What it means

Structured data measures whether the page contains machine-readable markup (in JSON-LD, Microdata, or RDFa format) that explicitly declares the type of content, the organisation behind it, and the specific entities involved. The three most valuable types for AEO are Organisation schema, Service or Product schema, and FAQPage or HowTo schema.

Evaluation criteria

Schema markup present

Any schema.org JSON-LD or Microdata block is present in the page source.

Organisation/Service schema

An Organisation or Service schema is present with name, URL, description, and sameAs links.

FAQPage/HowTo schema

A FAQPage or HowTo schema wraps question-and-answer content on the page.

Pass = schema present and well-formed. Warn = schema present but incomplete. Fail = no schema detected.

General recommendations

- Add a JSON-LD block in the for every key page type: WebPage, Organisation, and either FAQPage or HowTo.
- Include sameAs links to LinkedIn, Wikidata, and Crunchbase in Organisation schema to strengthen entity disambiguation.
- For service pages, add Service schema with serviceType, provider, and areaServed properties.
- Validate schema using Google's Rich Results Test before publishing.

Bad example

A service page with no schema markup at all. AI engines read the page text and attempt to infer what the organisation does, what services it offers, and what questions it answers — without any machine-readable confirmation. The result is low-confidence extraction that rarely surfaces in AI answer results.

Good example

A homepage with a fully populated Organisation schema (name, url, logo, description, sameAs: LinkedIn, Wikidata), a Service schema for each major offering (name, serviceType, provider, areaServed), and a FAQPage schema wrapping the six most common prospect questions. AI engines can read the page structure directly without relying on natural language inference.

D2 Content Clarity

Why it matters

AI engines generate answers by identifying the clearest, most direct passage in a page that addresses a user's query. Pages written in vague marketing language, or that bury the key point inside long paragraphs, are systematically skipped in favour of pages that answer clearly and concisely. Content clarity is the dimension most directly under the control of content writers.

What it means

Content clarity measures whether the page states its core value proposition in a single scannable sentence near the top, whether it uses direct answer structures (question followed immediately by a one-sentence answer), and whether its heading hierarchy is logical and keyword-rich enough for AI engines to navigate.

Evaluation criteria

Clear value proposition

The page states what it offers, for whom, and what outcome it delivers in the first visible paragraph.

Direct concise answers

Key claims and answers are stated directly, without preamble or qualification, in a format AI engines can extract as a standalone snippet.

Heading hierarchy

The page uses a logical H1 → H2 → H3 structure with keyword-rich headings that signal topic and subtopic to crawlers.

Pass = clearly present and prominent. Warn = present but buried or qualified. Fail = absent.

General recommendations

- Lead every page with a one-sentence value statement that names the product/service, the target audience, and the primary benefit.
- Rewrite FAQ answers to follow an inverted pyramid: direct answer first, then supporting detail.
- Ensure every H2 is phrased as a standalone statement or question that makes sense out of context.
- Avoid opening sentences with "We are", "Our", or "Welcome to" — AI engines prefer third-person entity descriptions.

Bad example

A homepage whose opening paragraph reads: "We are a full-service digital marketing agency committed to helping our clients achieve their goals through innovative, data-driven strategies tailored to their unique needs." No entity named, no audience defined, no specific outcome stated.

Good example

A homepage whose opening paragraph reads: "Kabloom Agency delivers congress marketing campaigns for professional associations and medical societies worldwide, with an average 47% uplift in delegate registrations within 12 months." Entity named, audience defined, outcome stated with a metric.

D Entity Definition

Why it matters

AI engines build a knowledge graph of named entities — organisations, people, places, products — and their relationships. When a page clearly defines the entity it represents, AI engines can confidently attribute information to it and surface it for queries that reference that entity. Poorly defined entities get merged with other organisations, misclassified, or simply excluded from the knowledge graph.

What it means

Entity definition measures whether the page unambiguously names the organisation, defines the industry and services it operates in, and provides enough geographic and contextual signals for AI engines to correctly classify and disambiguate it from similar entities.

Evaluation criteria

Brand entity clarity

The page names the organisation explicitly, describes what it does, and uses consistent branding throughout.

Services named precisely

Each service or product is named with a specific label, not a generic category description.

Industry and geo context

The page names the specific industry vertical and geographic markets served.

Pass = signal clearly present. Warn = signal present but vague. Fail = signal absent.

General recommendations

- Use the exact same organisation name, product names, and service labels across all pages — consistency is how AI engines build entity confidence.
- Name the specific industry verticals served (e.g. "professional associations", "medical congresses") rather than generic terms ("events industry").
- Include the country or region in body copy and in schema areaServed properties.
- Add a brief "About [Organisation]" paragraph on every page — not just the About page — to anchor entity signals throughout the site.

Bad example

A services page that refers to "our marketing solutions", "digital campaigns", and "event promotion" without naming the specific services, the industry, or the geographic market. AI engines cannot reliably classify this as distinct from any other digital marketing agency.

Good example

A services page that states: "Kabloom Agency delivers multi-channel congress marketing campaigns for professional medical associations and patient advocacy societies across Europe, North America, and the Middle East. Services include delegate acquisition campaigns, sponsor activation programmes, and membership revenue growth." Specific entity, specific services, specific geography.

D4 FAQ & Direct Answers

Why it matters

The primary use case for AI answer engines is answering questions. Pages that explicitly contain question-and-answer pairs — especially when marked up with FAQPage schema — are structurally pre-formatted for AI extraction. A page with five well-formed FAQ items has five potential AI citation opportunities that a page without FAQs simply cannot match.

What it means

This dimension measures whether the page contains a visible FAQ section with question-format headings, whether answers are direct and concise enough to be extracted as standalone snippets, and whether any statistics or claims are attributed to a named source.

Evaluation criteria

FAQ section present

The page contains at least one section with question-format headings followed by direct answers.

Question-format headings

At least two H2 or H3 headings are phrased as questions in natural language.

Stats substantiated

Numerical claims or statistics reference a named source or study.

Pass = clearly present. Warn = present but weak or unformatted. Fail = absent.

General recommendations

- Add a minimum of five FAQ items to every key page, targeting the exact questions prospects are most likely to ask AI assistants.
- Lead each answer with a single direct sentence — the extractable snippet — before adding supporting detail.
- Use H3 tags for FAQ questions and follow each immediately with an answer tag for clean DOM structure.
- Attribute statistics to a named source: "47% average uplift across Kabloom campaigns 2018–2024" is citable; "significant uplift" is not.

Bad example

A page with no FAQ section. Prospects and AI engines must infer answers from paragraph text, which is harder to extract and less likely to be cited.

Good example

A page with an H2 "Frequently Asked Questions" section containing five H3 questions, each followed by a two-sentence answer (direct response + one supporting detail), all wrapped in FAQPage JSON-LD schema. Every answer is a potential standalone AI citation.

P Technical Signals

Why it matters

Technical signals are the infrastructure layer that AI engines rely on to correctly categorise, attribute, and deduplicate pages. Poor technical hygiene does not prevent a page from being indexed, but it reduces the confidence with which AI engines attribute content — leading to lower citation rates even for pages with excellent content.

What it means

This dimension measures whether the page has a complete and keyword-rich meta description, correct canonical and Open Graph tags, and whether its key content is accessible to crawlers without requiring JavaScript execution.

Evaluation criteria

Meta description quality

The meta description is present, under 160 characters, and includes the primary keyword and entity name.

Canonical and OG tags

A canonical URL is set, and og:title, og:description, and og:image are all present.

JS-dependent content risk

Key content (value proposition, FAQ answers, entity description) is present in the raw HTML response, not injected by JavaScript after load.

Pass = correctly implemented. Warn = partially implemented or suboptimal. Fail = absent or broken.

General recommendations

- Write meta descriptions as a sentence that contains the entity name, the primary service, and one specific differentiator.
- Ensure og:title matches the H1 exactly, and og:description matches the meta description.
- Test the page with Google's Mobile-Friendly Test or a raw HTML fetch to confirm key content is server-rendered.
- For React/Vue/Angular pages, implement server-side rendering or static generation for all marketing content.

Bad example

A page where the FAQ answers, client testimonials, and About section are all rendered by JavaScript after page load. A crawler that does not execute JavaScript sees a near-empty page with just the navigation and a loading spinner — no content to index or cite.

Good example

A page where all marketing copy, FAQ answers, entity descriptions, and schema markup are present in the raw HTML response from the server. JavaScript is used only for interactive elements like the audit tool itself, not for the static content that AI engines need to read.

D 6 E-E-A-T

Why it matters

E-E-A-T (Experience, Expertise, Authoritativeness, Trustworthiness) is the framework AI engines use to assess whether a page comes from a credible, qualified source. A page can score well on all other dimensions and still be ignored if there are no signals confirming the content comes from someone with real expertise. For Kabloom and its clients, this is especially important in the credential-driven association and congress world — AI engines expect evidence of genuine sector authority.

What it means

E-E-A-T measures whether the page provides signals that allow AI engines to assess the credibility and authority of the content and its source. This includes author identification and credentials, trust indicators such as client lists and case studies, and content freshness — since outdated content is treated as less authoritative.

Evaluation criteria

Author and expertise signals

Named author or organisation with stated credentials, sector experience, or a byline linking to an author profile.

Trustworthiness indicators

At least one of: named client list, case study with results, industry award, press mention, or third-party endorsement visible on the page.

Content freshness

A visible publication or last-updated date within the past 12 months, or content referencing recent events or statistics.

Pass = signal clearly present and prominent. Warn = present but weak or buried. Fail = absent.

General recommendations

- Add a named author or contributor byline to every substantive page with a one-line credential.
- Include at least one trust signal per page: a client logo, a case study reference, or a stat with attribution.
- Add or update a "Last updated" date on key service and landing pages.
- Link author names to a short bio page or LinkedIn profile to create a named entity connection.

Bad example

A service page with no author name, no client references, no logos, no publication date, and no external validation of any kind. An AI engine has no basis to assess whether this is an authoritative source.

Good example

"Published by the Kabloom Agency team · Updated July 2026. Kabloom has delivered congress marketing campaigns for 100+ professional associations. Our average client sees a 47% uplift in registrations within 12 months (measured across campaigns 2018–2026)." Named organisation, specific credentials, substantiated claims, recency signal.

How Scores Are Calculated

Each of the six dimensions receives a score from 0 to 100 based on how many checks pass, weighted by severity:

Result pattern	Typical score range
All 3 checks pass	90–100
2 checks pass, 1 warn	65–80
2 checks pass, 1 fail	50–65
1 check passes, 2 warn	45–55
1 check passes, 1 warn, 1 fail	30–45
All checks warn	30–40
1 pass, 2 fails	20–30
All checks fail	0–20

The **Overall AEO score** is the mean of the six dimension scores.

Understanding Possible Inconsistencies

You may notice that auditing the same page twice produces different scores. This is expected behaviour, not a bug. There are three root causes:

1. The web search tool returns different content each run

The audit tool uses AI-powered web search to fetch page content. Each search may return a different cached snapshot of the page — sometimes richer, sometimes sparser — depending on what the search index has stored at that moment. A richer snapshot will detect more signals and score higher than a minimal one.

2. Claude's analysis involves judgment

Even with temperature set to zero (which we do to minimise this), large language models apply nuanced judgment when evaluating signals like content clarity or E-E-A-T. A phrase that reads as a pass in one run may read as a warn in another if the model weights the evidence slightly differently.

3. The parallel call architecture

The audit runs two simultaneous API calls — each covering 9 checks. Both calls independently fetch the page. If one call gets a richer fetch than the other, the checks it covers may score differently.

What this means in practice:

- Do not treat individual scores as exact measurements — treat them as directional signals.
- A dimension score can vary by ± 10 -15 points between runs.
- For reliable trending, compare averages across 2-3 runs rather than individual audits.
- Consistent fails across multiple runs are reliable indicators of genuine issues.
- Consistent passes across multiple runs confirm signals are genuinely in place.
- Compare before/after averages (2-3 runs each) to measure the impact of improvements.

The most reliable use of the tool: audit a page before making AEO improvements, make the changes, wait 24-48 hours for search indexes to update, then audit again. The difference between before and after scores (averaged across 2-3 runs each) is a reliable measure of improvement.

Priority Order for Fixes

When a page scores below 70, address fixes in this order for maximum impact:

Priority	Fix	Rationale
1	Add schema markup (Organisation + FAQPage)	Highest signal-to-effort ratio. One JSON-LD block can lift Structured Data score by 4
2	Add or rewrite FAQ section	FAQ content directly addresses how AI engines extract answers. 5 well-formed Q&A
3	Add author byline and last-updated date	E-E-A-T fixes are fast to implement and directly address AI engine trust signals.
4	Rewrite opening paragraph as a value statement	Content clarity improvements affect how extractable all other content is.
5	Entity definition in body copy	Medium effort, improves knowledge graph inclusion over time.
6	E-E-A-T signals	Add author bylines, client references, and last-updated dates — strong authority sign
7	Meta description and OG tags	Quick wins for Technical score. Usually fixable in a CMS in under 10 minutes.
8	Server-render key content	High effort but high impact for JavaScript-heavy pages. Required for reliable AI craw

Document maintained by Kabloom Agency. For questions or updates contact matias@kabloom-agency.com