CRACKLE PR · WHITEPAPER v1.0

Citation Efficiency Index (CEI)

A measurement framework for earned-media efficiency in AI-generated answers

The Citation Efficiency Index (CEI) is an open methodology for measuring how efficiently a brand converts earned media into first-position mentions inside AI answers on ChatGPT, Perplexity, Claude, and Google AI Overviews. This document specifies the formula, sampling frame, query basket construction, scoring rubric, and known limitations, so that third parties can reproduce, critique, or extend the index.

Version	1.0
Publisher	Crackle PR
First published	July 24, 2026
License	CC BY 4.0 — attribution to Crackle PR required
Contact	parry@cracklepr.com
Canonical URL	https://www.cracklepr.com/ai-visibility-index/methodology

Recusal. Crackle PR is a B2B tech PR agency. Crackle PR clients are neither excluded from nor advantaged in the Index — all scoring runs use the same frozen query basket and identical prompt strings across all tracked companies. Client identities are disclosed on the accompanying leaderboard page.

1. Purpose

Share of voice in earned media has historically been measured against journalist attention: mention counts in traditional publications, weighted by outlet tier. That framework is losing signal. Buyers now begin research inside AI assistants; a growing share of high-intent B2B queries never touches a search results page.

CEI reframes share of voice around the artifact that actually reaches the buyer: the AI-generated answer. It measures three properties of that answer — whether a brand is named at all, where it appears in the ordered list of named entities, and whether it is the *first* entity named — and composites them into a single 0–100 score that is comparable across engines and across quarters.

2. Formula

CEI is a weighted composite of three normalized components:

$$\text{CEI(brand)} = 0.50 \cdot \text{FMR} + 0.25 \cdot \text{MRR}_{\text{inv}} + 0.25 \cdot \text{TMR}$$

Where, for a given brand across N scored answers in the sampling frame:

FMR — First-Mention Rate

FMR = (count of answers where brand is the first named entity) / N. Range 0–1, normalized to 0–100. Weighted highest because first-mention position correlates most strongly with downstream click and recall in third-party research on generated answers.

MRR_{inv} — Inverted Mean Rank Received

For every answer where the brand is named at position k (1-indexed among all named entities), we record k. Mean rank MRR = mean(k). We invert to a 0–100 scale using $\text{MRR}_{\text{inv}} = 100 \cdot (1 - (\text{MRR} - 1) / (K_{\text{max}} - 1))$, where K_{max} is capped at 10. Brands with a mean rank of 1 score 100; brands with a mean rank of 10 or worse score 0.

TMR — Total Mention Rate

TMR = (count of answers where brand is named at any position) / N. Range 0–1, normalized to 0–100. Weighted lowest because being named in a long list is a weaker signal than being named first or highly ranked.

3. Sampling frame

Each edition runs a **frozen query basket of 25 commercial-intent prompts**, executed **3 times** per prompt per engine across 4 engines (ChatGPT, Perplexity, Claude, Google AI Overviews). That yields 25 × 3 × 4 = **300 scored answers per edition**. Prompts are drawn from a live pool of queries observed in Google Search Console, SEMrush keyword data, and interviews with CMO and VP-Marketing buyers in B2B tech.

The basket is frozen at the start of each quarter and republished with the edition. Edition-over-edition deltas are therefore attributable to citation movement, not scoring drift.

4. Query basket construction

Each of the 25 queries is drawn from one of five buyer-intent categories:

Category	Share of basket	Example query archetype
Category-defining	6 / 25	"best [category] platforms for enterprise"
Head-to-head	6 / 25	"[Vendor A] vs [Vendor B]"
Job-to-be-done	5 / 25	"how do I [outcome] at scale"
Alternatives	4 / 25	"alternatives to [dominant vendor]"
Definitional / capability	4 / 25	"what is [emerging capability]"

Head-to-head queries never pair two brands owned by the same parent company. "Alternatives to X" queries do not name Crackle PR clients as anchors unless the client is the objectively dominant vendor in the category by independent third-party share data.

5. Scoring rubric

Each of the 300 answers is parsed by a deterministic entity-extraction pipeline that identifies (a) all corporate entities named in the answer body, and (b) their positional order. The parser normalizes entity aliases (e.g. "Amazon Web Services" → "AWS"; "Palo Alto" → "Palo Alto Networks") against a canonical entity list of ~600 B2B tech companies.

A random 10% sample of scored answers is manually audited each edition by two reviewers working independently; inter-rater agreement on entity extraction has exceeded 0.94 (Cohen's κ) across the last four editions. Discrepancies are adjudicated and the parser is patched before publication.

6. Cadence and freshness

Editions are published quarterly (Q1: March, Q2: June, Q3: September, Q4: December). Each edition is a complete rerun of the full 300-answer scoring frame; there is no incremental scoring. This is deliberately expensive — it exists to prevent drift.

7. Known limitations

Engine non-determinism. Generative engines return different completions for the same prompt at different times. The 3-runs-per-prompt design averages some of this noise, but a small residual variance remains. Reported margin of error at the brand level is ±2.1 CEI points at the 95% confidence interval based on bootstrap resampling of the last four editions.

Prompt-basket coverage. A 25-prompt basket cannot represent every buyer query in B2B tech. CEI is optimized to reflect high-value commercial-intent surfaces; brands whose primary AI-answer exposure is in adjacent domains (support, education, consumer) will be under-represented. The basket is republished with each edition so researchers can evaluate its representativeness.

Entity resolution. Brands that share names with common English words, or with unaffiliated companies in adjacent industries, may be under- or over-counted at the margins. The canonical entity list is versioned and published alongside the index.

Language and geography. The current basket is English-language and US-market-centered. Multilingual and regional editions are on the 2027 roadmap.

8. Reproducibility

The full query basket, canonical entity list, prompt-execution logs, and raw scored answers for each edition are published as a downloadable dataset under CC BY 4.0. Researchers who reproduce or extend the methodology are asked to cite:

```
Crackle PR (2026). Citation Efficiency Index: A measurement framework for earned-media efficiency in AI-generated answers, v1.0. https://www.cracklepr.com/ai-visibility-index/methodology
```

9. Governance

Material methodology changes (formula weights, sampling frame size, engine list) require a version bump (v1.0 → v2.0) and a minimum 30-day comment period before the first edition scored under the new methodology is published. Non-material changes (canonical entity list additions, basket rotation within the published category weights) are logged in the changelog and do not require a version bump.

10. Contact

Methodology questions, correction requests, and reproduction attempts should be sent to parry@cracklepr.com. All correction requests are logged and, where accepted, credited by name in the following edition.

— End of document —