CLOUDFLARE

From ranking to recommended: Answer engine optimization (AEO)

Ashish Vinodkumar
Senior Machine Learning Engineer, ETI
Cloudflare

Will Rowe
Senior Systems Engineer, ETI
Cloudflare

Discovery is evolving: SEO > AEO > Agentic Internet

STAGE 01

SEO: Rank

- Optimize pages for search indexes and ranked links
- Primary signal: position, impression, click

STAGE 02

AEO: Be recommended

- Optimize evidence and clarity for synthesized answers
- Primary signal: mention, citation, prominence

STAGE 03

Agentic: Be chosen

- Make products and services usable by agents acting for people
- Primary signal: successful action and outcome

AI crawling is shifting from training to answering



+116%
OAI-SearchBot, OpenAI's search crawler

+32%
ChatGPT-User: visits when a person asks ChatGPT

+7%
GPTBot, OpenAI's training crawler

Source: Cloudflare Radar, May 2025 vs May 2026

Four market gaps create an opening at the edge

GAP 01

Prompt bias

- A prompt panel is a sample, not demand itself.
- Results vary by wording, model, place, and time.

GAP 02

Probe cost

- Every prompt, model, and repeat adds cost.
- Coverage must be broad enough to trust and bounded enough to run.

GAP 03

Action gap

- A visibility score says what happened.
- Teams still need to know whether bots can access content and what to change.

GAP 04

Missing ground truth

- Synthetic answers rarely show real site behavior.
- The edge sees crawls, blocks, and answer-engine referrals.

Cloudflare connects the full AEO loop

CLOUDFLARE

STEP 01

Measure answers

- Benchmark realistic customer questions across leading assistants
- Track mentions, citations, and share of model

STEP 02

Observe behavior

- See which AI crawlers access your content
- Track which answer engines send people back to your site

STEP 03

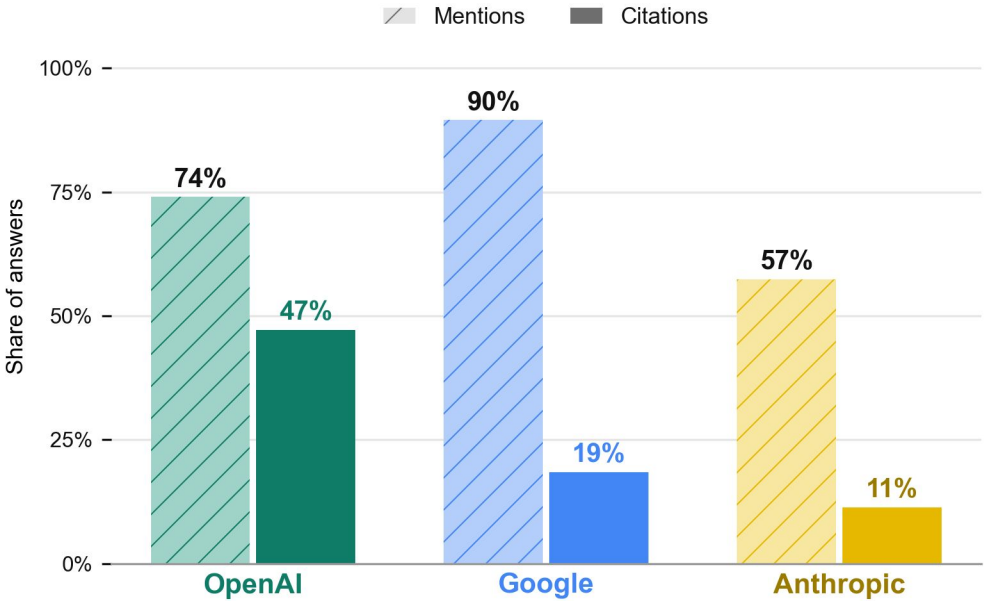
Take action

- Spot citation gaps and blocked retrieval bots
- Address unmanaged policies and execute next practical actions

Synthetic answers + edge signals + call to action = a more complete view of AEO

Does it matter which assistant you ask?

Yes. On the same questions, OpenAI linked to brands far more often, and Google named them without linking.



4×

OpenAI's citation rate was four times Anthropic's: 47% of answers vs 11%.

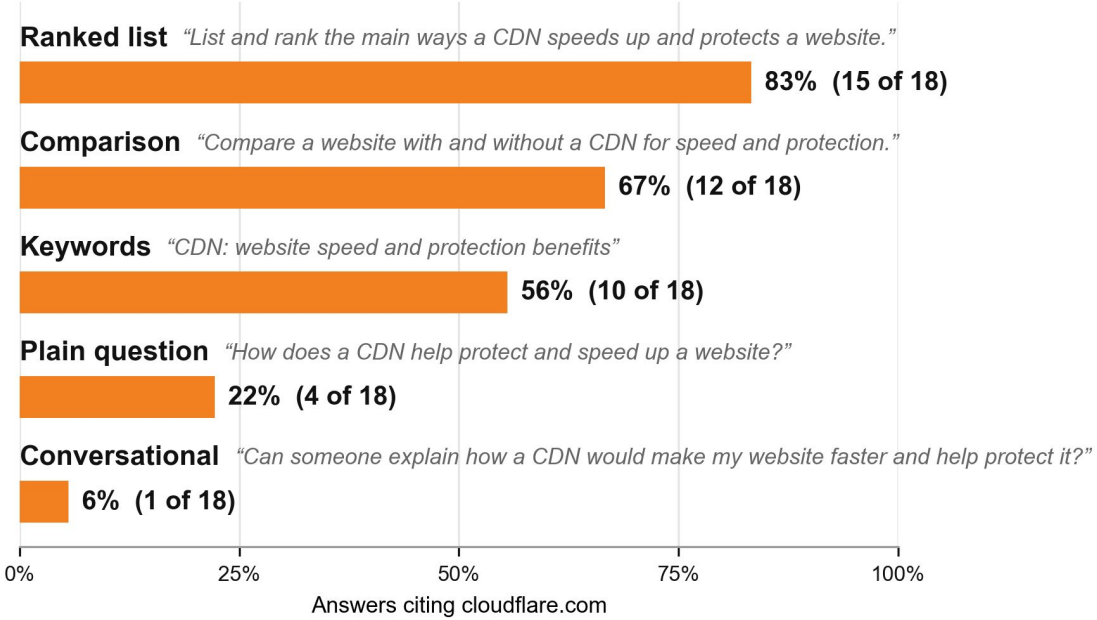
90% vs 19%

Google had the most mentions, but few citations.

Data: Cloudflare experiment, 22 Sep 2026. 50 questions on websites, hosting and security, 8 runs, 800 answers per assistant. OpenAI GPT-5.4 mini, Google Gemini 3.5 Flash-Lite and Anthropic Claude Haiku 4.5 APIs with web search, not the consumer apps. Citation = a link to a tracked brand's own site.

Does the wording change your score?

Yes. Asking for a comparison or a ranked list got brands cited more often. Here's one example.



31% vs 19%

of answers cited a brand when we asked for a comparison instead of a plain question, across 128 questions.

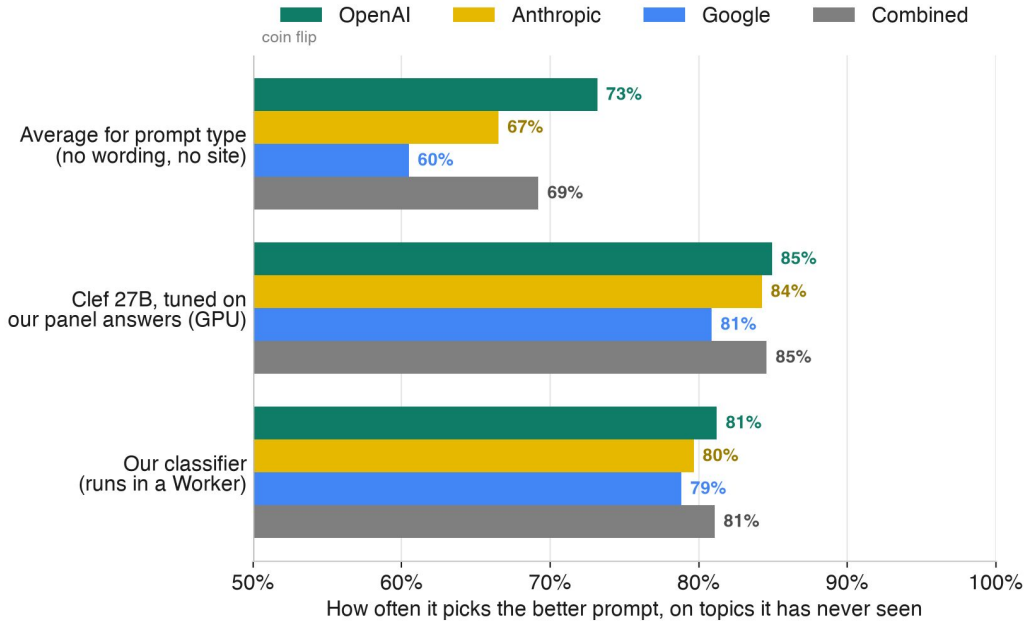
25% vs 19%

of answers cited a brand when we asked for a ranked list instead of a plain question.

Data: Cloudflare experiment, 13 Oct 2026. 128 questions (2 from each of 64 topics), 5 wordings each, 1,280 answers, one run. CDN example: 18 answers per wording, 3 runs, 21–22 Sep 2026. OpenAI GPT-5.4 mini and Anthropic Claude Haiku 4.5 APIs with web search, not the consumer apps.

Spend probes on prompts that can tell you something

Our prompt classifiers can screen AEO prompt before they cost probes or skew your score



81% accuracy

High predictive power (0.81 vs 0.70 baseline) even before a site has historical data.

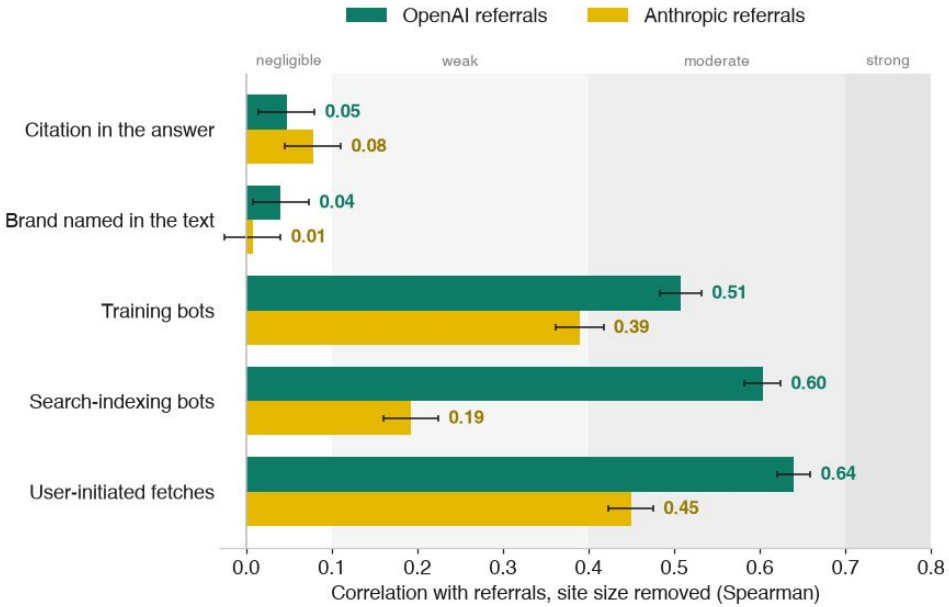
~10ms

A 3.6 MB distilled model running edge-side for near zero latency

Data: Cloudflare experiment, 13 Oct 2026. 314,751 assistant replies from 538 nightly panel runs, 16 topics, 23 Jul–9 Oct 2026. Claude Haiku 4.5, GPT-5.4 mini and Gemini 3.5 Flash-Lite APIs with web search, not the consumer apps. Experimental: missed its pre-registered bar on unseen topics by 0.005; per-brand ranking on new topics is still weak.

Signals from the network

AI traffic at the edge tracks real referral demand



0.45 - 0.64

User-initiated fetches track referrals most closely: a real person asked an assistant something that needed the site.

~0

Citations and brand mentions in prompt-panel answers barely track referrals once site size is removed.

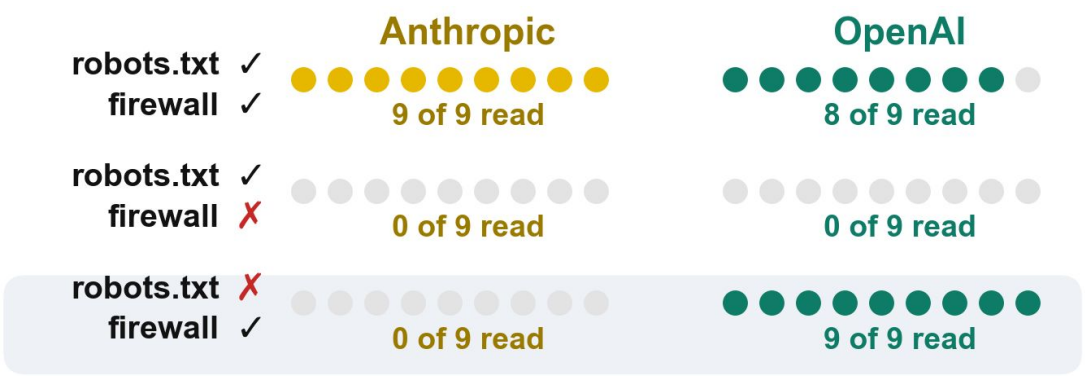
Data: Cloudflare experiment, 13 Oct 2026. Cloudflare network logs and panel runs, 23 Sep–8 Oct 2026, 3,521 sites. Partial Spearman correlation with AI referrals, site size (Radar popularity) removed; strength bands per Schober et al. 2018. Observational, not causal.

Does robots.txt keep AI bots out?

It depends on the bot. Only a firewall block kept both assistants out.

When a user shares a link

User bots (ChatGPT-User, Claude-User) · 27 test sites, link given



9 vs 0

sites read when robots.txt said no. OpenAI's user bot went in; Anthropic's stayed out, as each company's policy says.

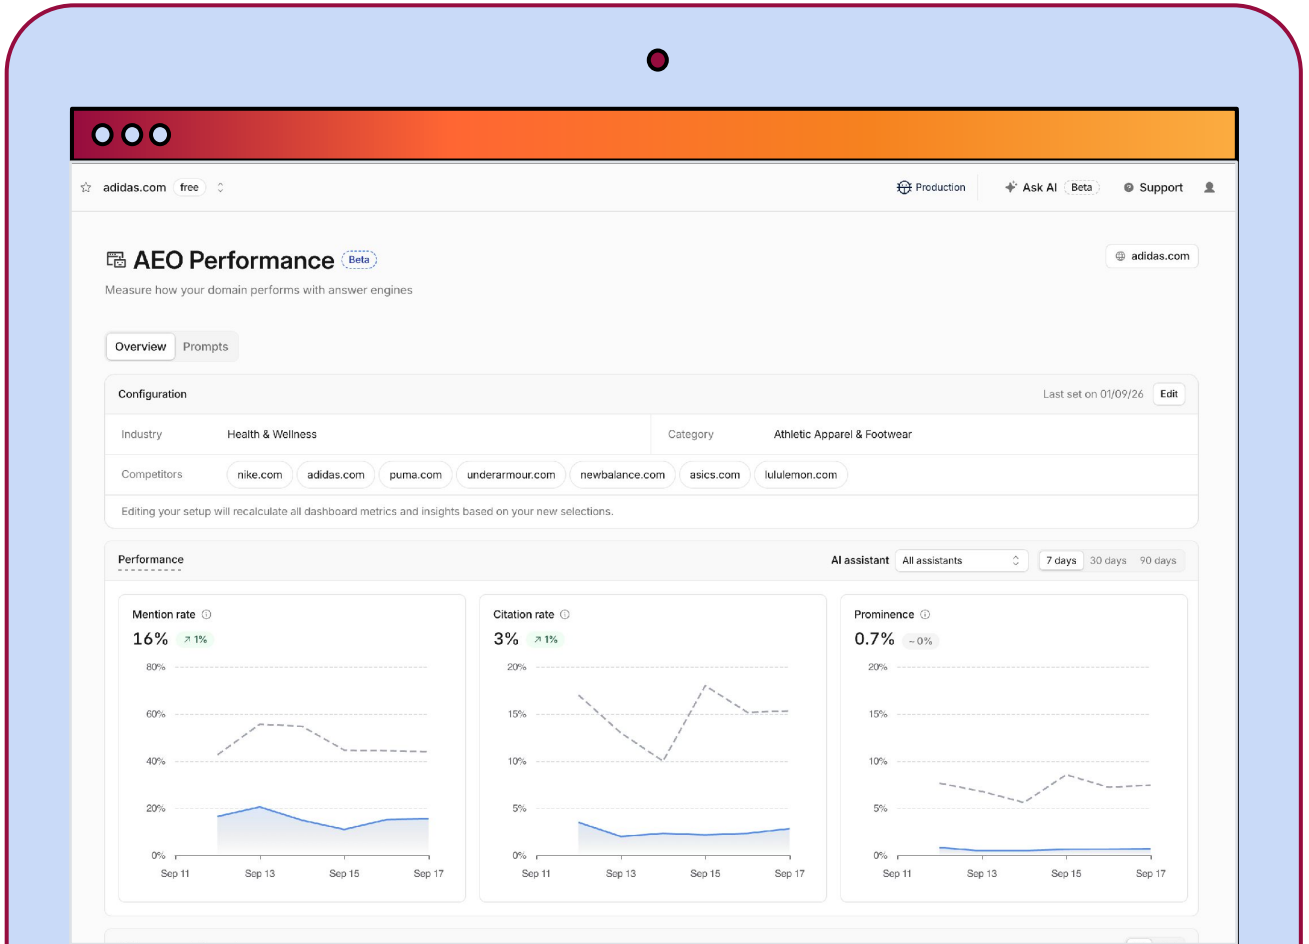
0 of 9

sites read by either assistant when the site returned a 403, even though robots.txt said yes.

Data: Cloudflare experiment, 15 Oct 2026. 27 test sites (9 per row) with made-up facts, 2 questions per site, each giving a page's link, 1 answer each; filled = the answer used the page's facts; the 403 went to both user bots. OpenAI GPT-5.4 mini and Anthropic Claude Haiku 4.5 APIs with web search, not the consumer apps. Policies linked: OpenAI and Anthropic crawler documentation.

Let’s see AEO in action

1. Choose your market
2. See recommendations and citations
3. Connect visibility to crawl and referral behavior



CLOUDFLARE

From ranking to recommended to chosen

Measure what AI says. Connect it to what the network sees.
Take action on access, content, and experience.

CLOUDFLARE

Your feedback is important

Please complete session survey in the Cloudflare Events app

CLOUDFLARE EVENTS

Open the app
↳ Click on the bell
↳ Tap “Survey”

