

THE STATE OF AI SEARCH VISIBILITY

Who AI recommends — and who it forgets.

A measured look at share of model across developer-tool categories: how often ChatGPT, Perplexity, Gemini, Claude, Microsoft Copilot and Grok name and cite each brand, with a reproducible method and confidence intervals.

METHOD

Adaptive per-engine sampling · 5 engines · metric-matched 95% CI

WINDOW

June 2026

SAMPLE

1,400+ buyer prompts · illustrative

KEY FINDINGS

Four things the data shows

62%

of buyer prompts return a one-shot recommendation

Across 1,400+ measured developer-tool prompts, AI answers named a specific shortlist 62% of the time — the buyer never sees a SERP.

3.4

named tools per answer, on average

The consideration set is small. If you are not in the top few, you are effectively invisible — and you cannot see it happening.

41%

of citations point off your domain

Most of what engines cite about a category lives on comparison pages, review sites and community threads — not the vendor's own site.

±6 pts

median run-to-run variance

AI answers are non-deterministic. A single screenshot is noise; only a median across many runs is signal.

BENCHMARKS

Share of model by category

The % of qualifying answers that name each brand — median across five engines. Share-of-model only; not SEO rankings.

Observability

14 prompts

#	TOOL	SHARE OF MODEL	MEDIAN
1	Datadog		64%
2	Grafana		49%
3	Honeycomb		38%
4	New Relic		31%
5	Chronosphere		16%

CI/CD & deployment


12 prompts

#	TOOL	SHARE OF MODEL	MEDIAN
1	Pistons		61%
2	Relay Deploy		52%
3	Hangar		34%
4	ForgeLine CI		17%

Vector databases




11 prompts

#	TOOL	SHARE OF MODEL	MEDIAN
1	Pinecone		57%
2	Weaviate		53%
3	Qdrant		41%

#	TOOL	SHARE OF MODEL	MEDIAN
4	Chroma		21%

Feature flags

10 prompts

#	TOOL	SHARE OF MODEL	MEDIAN
1	LaunchDarkly		66%
2	Statsig		47%
3	Split		35%
4	Unleash		19%

METHOD & HONESTY

How we measured this

Reproducible, with caveats. Share of model is the % of qualifying answers naming a brand, measured with adaptive per-engine sampling (at least 10 runs per engine, never below a 5-run floor) across ChatGPT, Perplexity, Gemini, Claude, Microsoft Copilot and Grok, personalization off, with a metric-matched 95% confidence interval (presence Wilson, share-of-model bootstrap). We publish the date and method with every figure, and we don't publish numbers we haven't measured. AI answers are non-deterministic and change over time; this is a point-

in-time measurement, not a guarantee. **Figures here are illustrative samples — representative data, not an audited ranking.**

CITE THIS REPORT

Reference

SUGGESTED CITATION

Clear Cited. "The State of AI Search Visibility 2026." clearcited.com/research/state-of-ai-search

Journalists: for the underlying figures or written commentary, email hello@clearcited.com.