



RARE BIRD LAB

FIELD GUIDE / GENERATIVE ENGINE OPTIMIZATION

GEO FOR DEVTOOLS

How to get your product cited by ChatGPT, Perplexity, and Google's AI Overview, when your buyer reads the answer and never scrolls to your link.

► 7 PLAYS + A ONE-PAGE CHECKLIST

THE FINISH LINE **MOVED**

For years the job was ranking a link and earning the click. Now a large share of technical buyers read the model's summary and stop there. If the summary skips you, most of that audience never learns you exist. GEO is how you get into the summary.

~15%

of Google searches showed an AI Overview by late 2025, after peaking near 25% mid-year.

SEARCH ENGINE LAND, 2025

8% → 1%

When an AI summary appears, people click a link 8% of the time, and a link inside the summary just 1%. Without a summary it is 15%.

PEW RESEARCH, 2025

84%

of developers now use or plan to use AI tools, and ChatGPT is the one most reach for first.

STACK OVERFLOW 2025 DEVELOPER SURVEY

780M

queries a month were flowing through Perplexity by mid-2025. Your buyer is asking it about your category.

TECHCRUNCH, 2025

GEO shares most of its fundamentals with SEO. Be genuinely useful, be crawlable, be the most credible source on the page. What changed is where the read ends: inside the answer itself, not on your site.

MAKE YOUR BEST CLAIM **EASY TO LIFT**

Seven moves that make a page quotable by an answer engine. The same things make a senior engineer trust it.

1 Answer the exact question in the first 100 words

Skimming engineers and answer engines both weight the opening heavily. Lead with the direct answer, then earn the rest of the read with real depth.

2 Give the model a line it can defend

Models reach for claims with a number and a source attached and skip the adjectives. A specific, sourced sentence is a liftable sentence.

DO "Cut build times 40% on a 500-repo benchmark."

NOT "Dramatically faster builds."

3 Write headings as the literal query

Use "how to connect X to Y" as a heading, not "Integrations." Match how the buyer phrased the search so both the reader and the model see the page answers it.

4 Structure for extraction

A TL;DR up top, one idea per section, comparison tables, and real code blocks. Give the engine clean, self-contained chunks it can quote without stitching.

EARN THE CITATION

5 Show working code, real config, actual output

A copy-pasteable snippet that runs is the strongest trust signal on a page. Screenshots of real output beat rendered mockups. If a claim can be demonstrated, demonstrate it.

6 Be specific, and honest about limits

Name the version, the constraint, the edge case, and where the tool is not the right fit. A page that admits limits reads as more credible, to engineers and to the models trained on what they trust.

7 Become the source that gets cited

Models cite the sources engineers already trust. Get quoted in docs, issues, and community answers using the exact phrasing people search. Human credibility and machine citations compound together.

When researchers at Princeton and Georgia Tech tested moves like these, their GEO methods raised page visibility in generative engines by up to 40%, with statistics, quotations, and citations to credible sources among the top-performing tactics.

SOURCE: AGGARWAL ET AL., ARXIV 2311.09735

THE ONE-PAGE CHECKLIST

Run any page you want cited through this before it ships.

- ☐ **Direct answer in the first 100 words.** A reader or a model gets the payoff before scrolling.
- ☐ **Every big claim has a number and a source.** No naked "faster," "simpler," "scalable."
- ☐ **Headings are written as real search queries.** Not "Features." Not "Overview."
- ☐ **A TL;DR, one idea per section, and a table** where a comparison belongs.
- ☐ **At least one snippet that actually runs** and one screenshot of real output.
- ☐ **You named a limit or a "not for you" case.** Credibility comes from the honest line.
- ☐ **The page is crawlable** and not buried behind a JS wall the engine cannot read.
- ☐ **You track a citation rate,** not just sessions. (See the next page.)

TRACK A CITATION RATE

Build a panel of the 40 prompts your buyers actually ask. Run them across ChatGPT, Perplexity, and Google's AI Overview on a schedule, and log who gets cited.

- If you run 40 prompts and get cited in 12, your citation rate is 30%. That number, tracked over time, is the metric worth reporting.
- Log which competitors show up across the same 40 answers and how often, then rank the whole set. Now you know where you actually stand.
- Getting cited in a third of answers is nothing to brag about when the leader hits 80%, but it is a solid position when the leader is stuck at 35%.

WANT YOUR PRODUCT IN THE ANSWER?

This is the work I do for devtools and AI startups: turning real technical depth into content that engineers trust and that answer engines cite back to you. If your category is getting answered by a model and it is not naming you, let's fix that.

► rarebirdlab.com/contact