

reaudit.

AI SEARCH VISIBILITY AUDIT

When a dev lead asks
an AI which tool
enforces
coding standards,
does
Vendor B show up?

A real-time, multi-engine audit of how Vendor B surfaces, and where it doesn't, across ChatGPT, Google AI, and Perplexity. Run today against the high-intent prompts engineering leaders are typing right now.

PERIOD

Audit run · June 3, 2026

PREPARED BY

Triantafyllos Rose Samaras
Reaudit · reaudit.io

01 · EXECUTIVE SUMMARY

You rank #2. But that's your **name** doing the work. On prompts that don't say "Vendor B," you appear **7%** of the time.

Across 21 AI answers to 6 high-intent prompts, Vendor B was named in **7** (33.3% overall), ranking #2 behind GitHub Copilot. But that number is almost entirely branded. On the 2 prompts that contain "Vendor B," you appear 100% of the time. On the 4 non-branded prompts a new buyer actually types, your visibility is **6.7%**, one mention out of fifteen. That gap is the entire opportunity.

OVERALL
VISIBILITY

33.3%

7 of 21 answers · #2
of 6 brandsNON-BRANDED
VISIBILITY

6.7%

1 of 15 category
answers · top of
funnel

AVG SENTIMENT

65/100

5 positive · 2 neutral ·
0 negative

CITATION RATE

100%

Every mention links
back to vendor-b.com

The three things that matter

1 · WHEN BUYERS KNOW YOUR NAME, THE ENGINES DELIVER, AND THEY LIKE YOU

On the 2 prompts that name you, Vendor B appears 100% of the time at consistently positive sentiment. The engines describe you as a "centralized rules engine," an "automated, centrally governed guardrail." The product story lands cleanly. The challenge is reach, not perception.

2 · THE CATEGORY BELONGS TO GITHUB COPILOT, YOU'RE NEARLY INVISIBLE TOP-OF-FUNNEL

On the 4 non-branded prompts a new buyer types, Vendor B appears once (ChatGPT only). The engines reach for **GitHub Copilot** by default. Three of the broadest questions, including "**enforce security policies in AI-generated code**," return zero Vendor B on every engine. 6.7% is the real number.

3 · AN ADJACENT DEV TOOL IS ALREADY BUYING CHATGPT ADS ON YOUR EXACT PROMPT

Sentry is running a ChatGPT sponsored card on the non-branded prompt "make AI coding assistants follow my company's coding standards," where Vendor B has zero organic presence. The paid layer of your category is open, and a non-tracked vendor moved first.

02 · THE BRAND WE'RE AUDITING

Vendor B, as you describe yourselves. And as the AI engines describe you.

Before we read the data, we read your site. This is the brand we calibrated the audit against, your positioning, your category, your stated ICP, and the language you've chosen to win with.

POSITIONING

Standards enforcement, before the first line is written. Vendor B gives every developer's AI agent your architecture rules, security policies, and coding standards automatically. Works with Cursor, Claude Code, and GitHub Copilot.

THE PROBLEM YOU SOLVE

AI-generated code drifts from a team's architecture and security rules. Vendor B acts as a centralized rule layer that auto-applies organization-wide standards so AI agents generate compliant code from the start, without manual review.

BRAND POINT OF VIEW

AI-generated code should never compromise on quality or security. Automate standards enforcement so developers focus on creativity and problem-solving, while AI handles the grunt work, safely. Quality assured by design, not by review.

MAIN KEYWORD

AI coding standards enforcement

"AI-generated code should never compromise on quality or security. Quality assured by design, not by review."

VENDOR B · BRAND POINT OF VIEW · VENDOR-B.COM

This audit asks one question of that brand: when a senior engineer or engineering manager at a 50 to 500 developer company asks ChatGPT, Google AI, or Perplexity for a way to enforce coding standards on AI-generated code, does Vendor B come up? And against whom?

IDEAL CUSTOMER (YOUR OWN DEFINITION)

A senior software engineer or engineering manager at a mid-to-large tech company (50 to 500 developers) responsible for code quality and standards. Pragmatic, efficiency-focused, fluent in AI coding tools, and frustrated by the inconsistency and manual effort of reviewing AI-generated code.

VOICE

Authoritative yet approachable, like a senior engineer mentoring a junior team. Direct, solution-oriented, confident, backed by real-world examples. Efficiency and best practices over theory.

COMPETITORS YOU NOMINATE

GitHub Copilot · CodeRabbit · Codex (OpenAI) · Tabnine · Sourcegraph Cody

03 · METHODOLOGY

What we tracked, how we tracked it, and why the numbers are real.

Generative engines are now the consideration-set layer above traditional search. We measured Vendor B's presence in that layer using a controlled, repeatable prompt set, not anecdotal queries, not screenshots, not vibes.

The prompt set

6 tracked prompts in your Vendor B onboarding set, each modeled on real engineering-buyer intent:

- How to make AI coding assistants follow my company's existing coding standards automatically
- Why does my AI agent write code that doesn't match our architecture rules?
- Tool to enforce security policies in AI-generated code before writing starts
- Automatically inject coding standards into Cursor, Claude Code, or Copilot
- Vendor B vs custom rules for AI coding assistants
- Is Vendor B secure and reliable for enforcing coding standards?

The engines

ChatGPT (OpenAI), **Google AI** (Search Generative Experience), and **Perplexity**, the three engines that intercept the overwhelming majority of consideration-stage research queries today.

The window

This audit was run today, against all three engines, from US-based residential infrastructure (your stated market). Every answer was parsed for brand mentions, cited sources, sentiment, and (for ChatGPT) sponsored ad cards. From here the same prompt set runs weekly inside Reaudit, so Vendor B sees movement the moment it happens.

What gets counted

A **mention** is an unambiguous brand name reference in the engine's answer. A **citation** is a clickable source link the engine attributes the answer to. **Sentiment** is scored on the sentence-level context where the brand appears, on a 0 to 100 scale.

Sample size

7 prompt-runs × 3 engines = **21 AI answers** analyzed today, across 5 tracked competitors and Vendor B. Citations harvested across the same surface: **230 total**, from **200 unique sources**.

WHY WE REPORT TWO VISIBILITY NUMBERS

Two of the six prompts contain the word "Vendor B." An engine almost always names a brand when the question already includes it, so branded prompts inflate a single visibility score. Throughout this report we separate **branded visibility** (the 2 prompts that name you) from **non-branded visibility** (the 4 category and problem prompts a new buyer types). The non-branded number is the real top-of-funnel signal, and for Vendor B the two are 15x apart.


WHY THIS MATTERS MORE THAN RANK-TRACKING

In a Google SERP, the buyer scrolls past your link or doesn't. In an AI answer, the engine has already made the shortlist for them. If Vendor B is not on the shortlist, the buyer never sees the option. The audit measures shortlist inclusion, the single most important new top-of-funnel metric in 2026.


04 · THE HEADLINE NUMBER

Vendor B appears in **33.3%** of answers. The number that matters is **6.7%**.

The 33.3% overall rate is real, and #2 is a strong place for a young brand. But it bundles two very different questions: prompts that name you, and prompts that don't. Split them, and the gap is stark.



Where the 7 mentions came from, by engine



Per-engine rate across all 6 prompts (branded included), captured in today's audit.

THE BRANDED PROMPTS ARE CARRYING THE ENTIRE SCORE

6 of your 7 mentions come from the two prompts that name you ("Vendor B vs custom rules," "is Vendor B secure"), which return you on every engine. Only one mention, ChatGPT on "inject standards into Cursor / Claude Code / Copilot," comes from a non-branded prompt. Strip the branded prompts and Vendor B is at 6.7%.

SENTIMENT IS SOLID, WITH ONE FIXABLE HEDGE

Average sentiment is 65 of 100. Where the engines describe you in depth they are warm: "centralized rules engine," "automated guardrail," "validation layer." The two neutral scores are ChatGPT's "promising, but not yet proven" caution and one list mention. You don't have a perception problem. You have a non-branded reach problem, which is the most fixable kind.

05 · COMPETITOR BENCHMARK


The real race is **you vs GitHub Copilot**. Three of your tracked rivals are invisible.

Same 6 prompts, same audit. This is the leaderboard.

	BRAND	DOMAIN	MENTIONS	CITATIONS	SENTIMENT	SHARE
1	GitHub Copilot	github.com	9	7	52	52.9%
2	Vendor B (you)	vendor-b.com	7	7	65	41.2%
3	CodeRabbit	coderabbit.ai	1	1	80	5.9%
4	Codex (OpenAI) · Tabnine · Sourcegraph Cody	3 brands	0	0	n/a	0%

Share = each brand's mentions as a percentage of the 17 total category mentions (all prompts, branded included).

Share of voice, visualized



YOUR #2 RANK IS REAL, BUT BRANDED-DRIVEN

Strip out the two prompts that name you and the picture flips: GitHub Copilot is the default the engines reach for on every non-branded, category-level prompt, while Vendor B appears once. The #2 finish comes from owning your own name, not the category. Closing that is recommendation #1.

THE DEDICATED-ENFORCEMENT NICHE IS WIDE OPEN

Your closest direct competitor, CodeRabbit, has just 1 mention. Copilot is a general assistant, not a standards-enforcement tool. Vendor B already leads the specific category of "enforce standards on AI-generated code" on coverage and citation rate. The job is to make the engines surface that category on non-branded prompts before anyone else claims it.

06 · PROMPT × ENGINE PERFORMANCE

You win every prompt that names you. You vanish on the ones that don't.

Grouping every tracked prompt by type makes the inflation impossible to miss: branded prompts are all green, the non-branded block is almost entirely grey.

PROMPT	CHATGPT	GOOGLE AI	PERPLEXITY
BRANDED PROMPTS · CONTAIN YOUR NAME			
Vendor B vs custom rules for AI coding assistants	✓	✓	✓
Is Vendor B secure and reliable for enforcing standards?	✓	✓	✓
NON-BRANDED PROMPTS · TOP OF FUNNEL			
Automatically inject coding standards into Cursor, Claude Code, or Copilot	✓	.	.
Make AI assistants follow my company's coding standards	.	.	.
Why does my AI agent ignore our architecture rules?	.	.	.
Tool to enforce security policies in AI-generated code	.	.	.

Green ✓ = Vendor B named in that engine's answer. Light green = one engine only. Grey "." = absent.

YOUR ONE NON-BRANDED Foothold IS THE "INJECT STANDARDS" PROMPT

"Automatically inject coding standards into Cursor, Claude Code, or Copilot" is the only non-branded prompt where Vendor B surfaces, and only on ChatGPT, listed among "emerging approaches." It maps exactly to your tagline. Winning Google AI and Perplexity on this same prompt is the nearest, highest-probability gain.

THREE CORE PROMPTS RETURN ZERO, AND THEY ARE YOUR CATEGORY

"Make AI assistants follow coding standards," "why does my AI agent ignore our architecture rules," and "enforce security policies in AI-generated code" return no Vendor B on any engine. These are the broadest, earliest questions your buyer asks, and they describe your product almost word for word.

07 · VERBATIM · HOW AI TALKS ABOUT YOU

The actual sentences. Unedited. From the engines.

All 7 mentions, with the prompt that triggered them and the language the engine used. This is the qualitative ground truth behind the visibility score, including the one caution worth fixing.

CHATGPT

Sentiment 60 · positive · 14 citations

"Vendor B vs custom rules for AI coding assistants"

The key difference is that **it tries to act as a centralized rule-management and validation layer for coding agents**, whereas custom rules are usually just files like AGENTS.md, CLAUDE.md, or .cursor/rules.

GOOGLE AI

Sentiment 60 · positive · 14 citations

"Vendor B vs custom rules for AI coding assistants"

You are choosing between

an automated, centrally governed guardrail solution and a manual, file-based context injection system. Both approaches solve the issue of "architectural drift," where AI assistants deviate from a team's rules.

PERPLEXITY

Sentiment 70 · positive · 10 citations

"Is Vendor B secure and reliable for enforcing coding standards?"

Vendor B appears to be a centralized rules engine and validation layer for AI coding tools that aims to enforce coding standards, security, and architectural guidelines before code is generated.

CHATGPT

Sentiment 55 · neutral · 12 citations

"Is Vendor B secure and reliable for enforcing coding standards?"

I would describe Vendor B as **promising, but not yet proven** at the level where I'd blindly trust it as the sole enforcement mechanism. Vendor B's core idea is to centralize engineering rules (coding standards, security).

CHATGPT

Sentiment 50 · neutral · 15 citations · only non-branded mention

"Automatically inject coding standards into Cursor, Claude Code, or Copilot"

There are a few emerging approaches. **Centralized standards plus automatic sync** is the most common: instead of maintaining separate CLAUDE.md, .cursor/rules, and copilot-instructions.md files, you keep one source of truth.


THE ONE CAUTION TO ENGINEER AWAY

ChatGPT's **"promising, but not yet proven"** is the single hedge in an otherwise glowing set. It is a trust-signal gap, not a product gap. Third-party proof (reviews, security posture page, case studies) is what turns that sentence into an unqualified recommendation.

08 · THE SOURCE GRAPH

You're in the directories. You're missing from the **authorities** that answer your category.

AI engines synthesize from a small set of trusted sources. Vendor B already shows up in product directories and an early review. The gap is the security and architecture authorities the engines cite on your non-branded prompts.



The third-party sources surfacing your category

Top block: sources that already mention you. Bottom block: authorities the engines cite that do not.

<p>vendor-b.com</p> <p>Your own domain · the single most-cited source in the set · ChatGPT, Google AI & Perplexity</p>	17 citations
<p>producthunt.com / vendor-b</p> <p>Product directory · names Vendor B · Google AI & Perplexity</p>	4 citations
<p>automated.com · everydev.ai · docs.vendor-b.com</p> <p>Reviews and dev-tool directories · name Vendor B</p>	3 each
<p>reddit.com</p> <p>Community threads on AI agents drifting from team rules · no Vendor B</p>	6 citations
<p>best.openssf.org · snyk.io · gitguardian · blackduck · harness</p> <p>Security authorities cited on your security prompt · no Vendor B</p>	2 each

THE EARLY FOOTPRINT IS REAL, KEEP FEEDING IT

ProductHunt, everydev.ai, productdirs, and the automated review already name Vendor B and feed Perplexity and Google. For a brand this young, being in the directory layer at all is a head start. Keep these listings fresh and detailed.

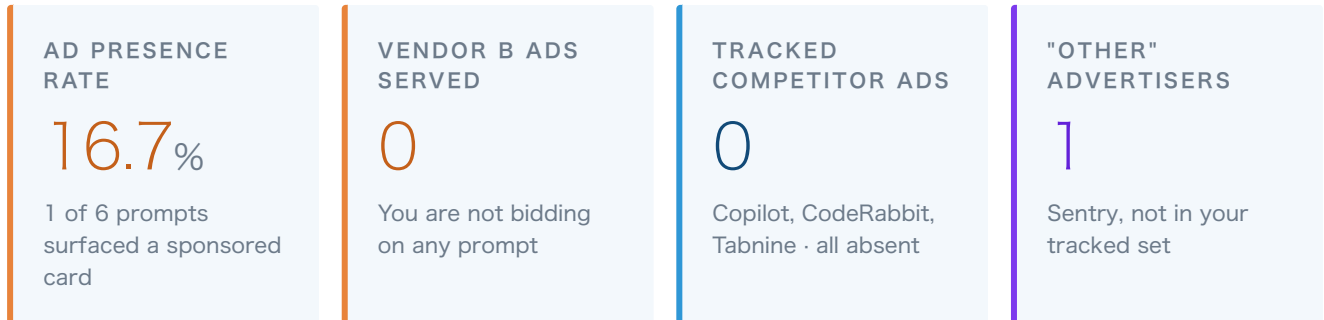
THE OPENING IS THE AUTHORITY LAYER

On the security prompt the engines cite OpenSSF, Snyk, GitGuardian, BlackDuck, and Harness, none of which mention Vendor B. That is why that prompt returns zero. The Reddit discussions Google cites on AI agents drifting from team rules are your exact use case. Getting named in these is the highest-leverage move for the prompts you miss.

09 · CHATGPT SPONSORED ADS · THE OPEN LANE

One advertiser is in your category's ChatGPT cards. It isn't you, and it isn't a competitor you track.

OpenAI's sponsored carousel is becoming the default monetization layer on ChatGPT. Across your tracked prompts we captured the paid cards. The channel is almost empty, which is both the opportunity and the warning.



Who's buying ChatGPT ads on Vendor B's prompts

ADVERTISER	DOMAIN	CARDS	PROMPTS HIT	CLASSIFICATION
Sentry	sentry.io	1	1	Adjacent: AI observability & error monitoring moving up the dev workflow

THE FIRST-MOVER WINDOW

With one advertiser and zero competitor presence, a modest ChatGPT ad budget on your six prompts buys near-uncontested share of the paid layer. Pair it with the non-branded prompts you miss organically and you cover both surfaces at once, before GitHub Copilot or CodeRabbit notice the lane is open.

WHY IT STILL MATTERS THAT IT'S NEARLY EMPTY

Ad inventory in dev-tools is wide open today. It will not stay that way. The same dynamic in adjacent B2B SaaS categories went from "no ads" to "every high-intent prompt has a card" in under two quarters. The audit will catch the shift the week it starts.

Read the one card buyers already see **before** they see your name.

The sponsored ChatGPT carousel card captured against your tracked prompts in today's audit.

Sentry

sentry.io

AI observability for AI teams

Get AI inference monitoring, error tracking, and performance insights.

Served on: "How to make AI coding assistants follow my company's existing coding standards automatically"

THE COMPETITIVE READ

Sentry is not a standards-enforcement tool. It sells AI observability and error tracking. But it is paying to appear on the single most valuable non-branded prompt in your category, the same prompt where Vendor B has zero organic presence. That is the pattern to watch: adjacent dev tools using paid ChatGPT cards to intercept buyers before they ever reach a shortlist.

THE ASYMMETRY IN YOUR FAVOR

None of your tracked competitors are bidding yet, and none of them are purpose-built for standards enforcement the way you are. A small, surgical ad spend on the "follow our coding standards" prompt would let you claim the paid layer on the exact query that describes your product, while it's empty.

CAVEAT WE'LL STATE PLAINLY

This is a same-day snapshot, so the ad layer is thin by definition. Treat the 16.7% presence rate as a baseline, not a trend. The value is in watching it move week over week, which the standing audit does automatically.

10 · THE PLAYBOOK

Seven moves to turn a branded #2 into category ownership.

Every recommendation maps to a finding earlier in this deck. No generic SEO. Each item is a leverage point in the source graph that drives AI engine answers, aimed squarely at the 6.7% non-branded gap.

- 1 Win the non-branded prompts you currently miss**

"Make AI assistants follow coding standards," "why does my AI agent ignore our architecture rules," and "enforce security policies in AI-generated code" return zero Vendor B. Publish citable assets in that exact language. This is the number that moves your real visibility from 6.7%.
- 2 Get named by the security and architecture authorities**

OpenSSF, Snyk, GitGuardian, BlackDuck, and Harness are cited on your security prompt and never mention you. Contribute guidance, get listed in their tool round-ups, and publish a piece that earns a citation alongside them.
- 3 Own the "architectural drift" narrative**

Google already cites Reddit discussions on AI agents drifting from team rules. That is your exact problem statement. Be the answer in those threads and publish the definitive "architectural drift" explainer.
- 4 Build third-party proof to kill the "not yet proven" hedge**

ChatGPT called you "promising, but not yet proven." Ship G2 reviews, named case studies, and a public security-posture page. That converts the one cautious sentence into an unqualified recommendation.
- 5 Widen the branded wedge into comparisons**

You win "Vendor B vs custom rules" and "is Vendor B secure" on all three engines. Publish your own "Vendor B vs AGENTS.md / .cursor/rules" comparison and a reliability page, so the surfaces the engines cite stay yours.
- 6 Take the open ChatGPT ad inventory**

One adjacent vendor (Sentry) is bidding; no competitor is. Run modest sponsored cards on your six prompts, prioritizing the "follow our coding standards" prompt, to cover the paid layer while it's uncontested.
- 7 Run the audit monthly, and track both metrics**

Watch overall AND non-branded visibility every month. The branded number will stay near 100%; the non-branded number is the one that signals whether you're catching GitHub Copilot on the category prompts that win new buyers.

From a one-time audit to a **continuous AI visibility practice.**

Everything in this document was generated by a single Reaudit project, run today. The same instrumentation runs every week, across every engine, every prompt, every competitor, and every sponsored ad. The findings become a live dashboard, with delta alerts when a competitor enters the source graph or a new advertiser bids against your prompts.

DATA ACCESS · REAUDIT MCP

Everything you just read is live, queryable data, exposed through the Reaudit MCP. So is more: every brand mention, every citation, every sponsored ad creative, every competitor move, every visibility gap, refreshed daily. Connect it to Claude, Cursor, or any LLM your team already uses, and Vendor B's founders pull this intelligence directly. No dashboards. No exports. Just ask.

STEP 1 · NOW

Workshop the playbook

60-minute working session with Vendor B's founders and growth lead to prioritize the 7 moves, assign owners, and lock the non-branded prompt push.

STEP 2 · 30 DAYS

Expand the tracked set

Scale from 6 to 40+ prompts across standards enforcement, architectural drift, AI code security, and Cursor / Claude Code / Copilot workflows.

STEP 3 · ONGOING

Monthly visibility review

Standing report on overall and non-branded visibility, with movement analysis, new ad threats, and source-graph recommendations, the first Monday of every month.

PREPARED BY

Triantafyllos Rose Samaras

Reaudit · AI Search Visibility & Brand Monitoring · reaudit.io

DISCUSS THIS REPORT

Book a working session →

reaudit.io · contact