

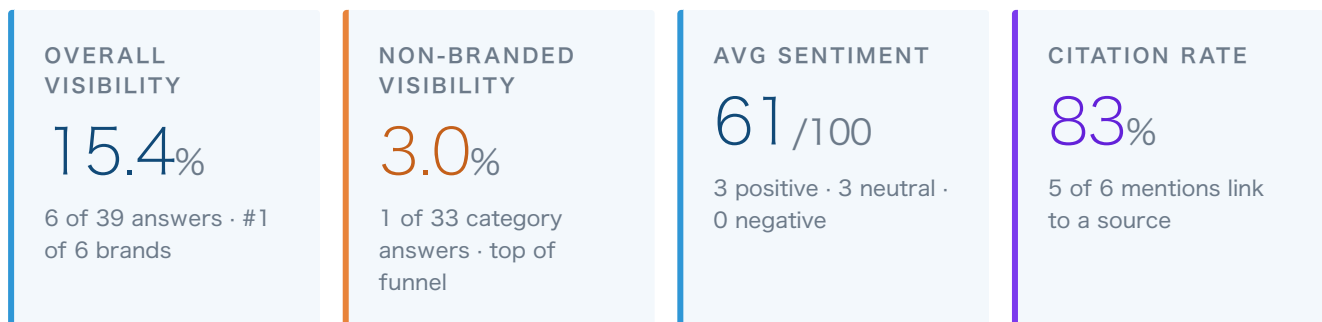
When an SRE asks an AI which platform catches incidents **before they ship**, does Vendor C show up?

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

01 · EXECUTIVE SUMMARY

You rank #1. It rests on **six mentions**, and on the questions an SRE actually asks, you appear **3%** of the time.

Across 39 AI answers to 16 high-intent prompts, Vendor C was named in **6** (15.4% overall), edging Datadog for #1 on a sentiment tiebreak. But that #1 is fragile: five of the six mentions sit on prompts that already say "Vendor C." On the non-branded prompts an SRE types, you appear **3.0%** of the time, once in thirty-three answers. And the engines aren't sure which Vendor C you are.



The three things that matter

1 · THE #1 RANK IS REAL, BUT IT IS A TIE ON A TINY BASE

Vendor C and Datadog each landed 6 mentions; you take #1 only on sentiment (61 vs 50). Six mentions is not category leadership, it is a toehold. Everywhere outside the prompts that name you, Datadog is the engines' default.

2 · TOP OF FUNNEL IS EFFECTIVELY ZERO: 3% NON-BRANDED

The SRE questions a buyer actually types, root cause analysis, MTTR, blast radius, alert fatigue, return Vendor C once across 33 answers. The broad reliability and AIOps prompts default to Datadog and generic tooling. This 3% is the real number, and the entire opportunity.

3 · THE ENGINES DON'T KNOW WHICH "VENDOR C" YOU ARE

ChatGPT openly asks "There appear to be two different companies using the Vendor C name. Which one do you mean?" Your source graph is diluted by **vendor-c.net**, **vendor-c.fi**, and unrelated namesakes. Disambiguation is a prerequisite for everything else.

02 · THE BRAND WE'RE AUDITING

Vendor C, 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

The Context & Control Model for Production.

Give production AI operational clarity. Built for the platform and SRE teams who own it, recognized in the Gartner Market Guide, 2026.

WHAT IT DOES

Context: catch the change that would have caused the incident before it deploys, with blast radius and a fix in hand. **Memory:** the next incident starts with everything your team has already solved. **Control:** runtime policy checks every agent action before it executes, and blocks anything that violates it.

SECURITY AND PRIVACY

SOC 2 Type II, GDPR, ISO 27001 in progress. SaaS or fully in-VPC. SAML SSO, SCIM, RBAC bound to your IdP. PII redaction before any model call. Bring your own model: Bedrock, Azure OpenAI, Vertex.

MAIN KEYWORD

AI SRE & production control

"AI agents need a proper model of production and a runtime that gates every action against it. We built both, as one system."

VENDOR C · FOUNDER POINT OF VIEW · VENDOR-C.AI

This audit asks one question of that brand: when a platform or SRE lead at a cloud-native company asks ChatGPT, Google AI, or Perplexity for the kind of incident, reliability, and production-control platform Vendor C claims to be, does Vendor C come up? And against whom?

IDEAL CUSTOMER (YOUR OWN DEFINITION)

Platform engineering and SRE teams who own production reliability at fast-moving, cloud-native companies. Responsible for incident response, deployment risk, and reliability, and tired of manually correlating logs and metrics to find root cause.

VOICE

Engineer to engineer. Calm, precise, operator-credible, backed by nine years building and operating observability at global scale. Ownership stays with engineers; the model is the control plane.

COMPETITORS YOU NOMINATE

Datadog · New Relic · Splunk · Honeycomb · Arize AI

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 C's presence in that layer using a controlled, repeatable prompt set, not anecdotal queries, not screenshots, not vibes.

The prompt set

16 tracked prompts across two sets, modeled on real SRE and platform-buyer intent:

- Best AI tool for automated root cause analysis in Kubernetes production
- How to reduce MTTR for cloud-native incidents using AI
- AI platform for blast radius analysis before code reaches production
- Best way to catch configuration drift before deployment
- Runtime policy enforcement for autonomous systems
- Vendor C reviews and case studies · Vendor C vs other platforms

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 C 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 by an LLM on the context where the brand appears, on a 0 to 100 scale with a neutral band.

Sample size

39 AI answers analyzed today across 16 tracked prompts and 3 engines, against 5 tracked competitors and Vendor C. Citations harvested across the same surface: **436 total**, from **407 unique sources**.

WHY WE REPORT TWO VISIBILITY NUMBERS

Two of the prompts contain the word "Vendor C." 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 category and problem prompts an SRE types). The non-branded number is the real top-of-funnel signal, and for Vendor C the two are roughly 28x 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 C 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 C appears in **15.4%** of answers. On the questions that don't name you, it's **3%**.

The 15.4% overall rate is real, and #1 reads well on a slide. But it bundles two very different questions: prompts that name you, and prompts that don't. Split them, and the picture changes completely.



Where the 6 mentions came from, by engine



Count of Vendor C mentions per engine across all tracked prompts, captured in today's audit.

THE BRANDED PROMPTS ARE CARRYING THE ENTIRE SCORE

Five of the six mentions come from the two prompts that name you ("Vendor C reviews and case studies," "Vendor C vs other platforms"). Only one mention, Perplexity on "tools to stop alert fatigue and automate incident triage," comes from a non-branded prompt. Strip the branded prompts and Vendor C is at 3%.

SENTIMENT IS FINE WHERE YOU ARE DISCUSSED

Average sentiment is 61 of 100. Where the engines describe you in depth they are positive: Google calls you "an incident management and automated root-cause analysis platform built for engineering teams." The neutral scores are the answers where Vendor C is one option in a list, or where the engine hedges on which Vendor C it means. Perception is not the problem. Reach and clarity are.

05 · COMPETITOR BENCHMARK


You edge **Datadog** for #1, on six mentions and a sentiment tiebreak.

Same prompts, same audit. This is the leaderboard.

	BRAND	DOMAIN	MENTIONS	CITATIONS	SENTIMENT	SHARE
1	Vendor C (you)	vendor-c.ai	6	5	61	37.5%
2	Datadog	datadoghq.com	6	5	50	37.5%
3	New Relic	newrelic.com	2	1	50	12.5%
4	Splunk	splunk.com	2	1	50	12.5%
5	Honeycomb · Arize AI	2 brands	0	0	n/a	0%

Share = each brand's mentions as a percentage of the 16 total category mentions (all prompts, branded included).

Share of voice, visualized



READ THE #1 HONESTLY: IT IS A COIN-FLIP WITH DATADOG

You and Datadog are level on mentions and citations. The tiebreak is sentiment, where the engines speak about Vendor C in your own product language and about Datadog generically. That is a genuine asset, but it sits on six data points. One week of Datadog content could flip the order.

THE CATEGORY DEFAULT, NOT THE LEADER, IS THE REAL OPPONENT

On the non-branded SRE prompts, the engines reach for Datadog and generic AIOps tooling by default, not for Vendor C. New Relic, Splunk, Honeycomb, and Arize barely register. The opening is not to out-mention Datadog on branded prompts, it is to own the problem prompts where no incumbent dominates yet.

06 · PROMPT × ENGINE PERFORMANCE

You surface on the prompts that name you. The category prompts are a wall of grey.

Grouping tracked prompts by type makes the inflation impossible to miss: the branded block is green, the non-branded block, the questions buyers actually type, is almost entirely absent.

PROMPT	CHATGPT	GOOGLE AI	PERPLEXITY
BRANDED PROMPTS · CONTAIN YOUR NAME			
Vendor C reviews and customer case studies	✓	✓	✓
Vendor C vs other AI governance platforms	✓	.	✓
NON-BRANDED PROMPTS · TOP OF FUNNEL			
Tools to stop alert fatigue and automate incident triage	.	.	✓
Automated root cause analysis in Kubernetes production	.	.	.
Reduce MTTR for cloud-native incidents using AI	.	.	.
Catch configuration drift before deployment	.	.	.

Green ✓ = Vendor C named in that engine's answer. Light green = one engine only. Grey "." = absent. Non-branded block shows 4 representative prompts of the larger category set.

YOUR ONE NON-BRANDED FOOTHOLD IS ALERT FATIGUE AND TRIAGE

"Tools to stop alert fatigue and automate incident triage for SREs" is the only non-branded prompt where Vendor C surfaces, and only on Perplexity, listed among options. It maps directly to your incident-resolution story. Winning ChatGPT and Google on this same prompt is the nearest, highest-probability gain.

YOUR CORE CATEGORY PROMPTS RETURN ZERO, AND THEY DESCRIBE YOUR PRODUCT

"Automated root cause analysis," "reduce MTTR with AI," "blast radius analysis," and "catch config drift before deployment" return no Vendor C on any engine. These are the broadest, earliest questions your buyer asks, and they restate Context, Memory, and Control almost word for word.

07 · VERBATIM · HOW AI TALKS ABOUT YOU

The actual sentences. Unedited. From the engines.

All 6 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 disambiguation issue worth fixing first.

GOOGLE AI

Sentiment 80 · positive · 20 citations

"Vendor C reviews and customer case studies"

Vendor C is an incident management and automated root-cause analysis platform built for engineering teams to diagnose production issues faster. Instead of manually correlating logs and metrics, the AI maps the system's cause-and-effect relationships to pinpoint the exact commit.

PERPLEXITY

Sentiment 75 · positive · 10 citations

"Vendor C reviews and customer case studies"

Vendor C is positioned as an incident resolution and alert triage platform that uses AI to automate investigations, reduce false positives, and accelerate root-cause analysis for engineering teams.

PERPLEXITY

Sentiment 60 · positive · 10 citations

"Vendor C vs other AI governance platforms"

Vendor C is an AI SRE platform focused on proactive reliability and incident speed, while other AI governance platforms vary in scope across policy, risk management, and model governance.

CHATGPT

Sentiment 50 · neutral · 12 citations · the disambiguation

"Vendor C reviews and customer case studies"

"There appear to be two different companies using the Vendor C name. Which one do you mean?" 1. Vendor C (SRE / reliability engineering platform)..." It then cites a customer case study, but only after asking the buyer to disambiguate.

PERPLEXITY

Sentiment 50 · neutral · 10 citations · only non-branded mention

"Tools to stop alert fatigue and automate incident triage for SREs"

A practical guide to reducing alert noise and accelerating triage, with Vendor C listed among the recommended approaches rather than led with. **One option among several**, not the default answer.

THE ONE ISSUE TO ENGINEER AWAY FIRST

ChatGPT's **"which Vendor C do you mean?"** is not a perception problem, it is an entity problem. The engines cannot cleanly resolve "Vendor C" to your company. Structured data, consistent Vendor C naming, and a clean entity footprint convert that hesitation into a confident answer.

08 · THE SOURCE GRAPH

You own vendor-c.ai. The gap is the **SRE authorities**, and the namesakes diluting your name.

AI engines synthesize from a small set of trusted sources. vendor-c.ai is already your most-cited domain. Two problems sit on top of it: the AIOps and SRE round-ups the engines cite on your non-branded prompts don't mention you, and unrelated "Vendor C" domains are polluting your entity.



The sources surfacing your category

Top block: your own footprint. Middle: gatekeepers the engines cite that do not name you. Bottom: namesakes diluting your entity.

vendor-c.ai

Your own domain · the single most-cited source in the set · Perplexity, Google AI & ChatGPT

18 citations

linkedin.com · featuredcustomers.com · incident.io

Proof and incident-tooling gatekeepers on your category prompts

2-4 each

reddit.com · cloudnativenow.com · g2.com

SRE communities and AIOps round-ups · "preventing production AI agents from going rogue" · no Vendor C

2-3 each

vendor-c.net · vendor-c.fi · unrelated-namesake.com

Unrelated namesakes sharing the name, diluting your entity

namesakes

THE OWNED FOOTPRINT IS REAL, KEEP FEEDING IT

vendor-c.ai is cited 18 times, more than any third-party source, and your customers already feed the engines through case studies and FeaturedCustomers. For a young brand, owning your own source graph is a head start most never get. Keep those pages fresh and schema-marked.

TWO OPENINGS: THE AUTHORITIES, AND THE NAMESAKES

On the SRE prompts the engines cite Reddit threads, CloudNativeNow, G2, and incident.io, none of which name Vendor C. Getting cited there is the highest-leverage organic move. Separately, vendor-c.net, vendor-c.fi, and unrelated namesake domains are confusing the engines about who you are. Both have to be addressed for non-branded visibility to move.

09 · CHATGPT SPONSORED ADS · THE OPEN LANE

Two advertisers are in your category's ChatGPT cards. Neither is you, and neither is 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 nearly empty of direct rivals, which is both the opportunity and the warning.

<p>PROMPTS WITH A CARD</p> <p>2</p> <p>both on non-branded SRE prompts</p>	<p>VENDOR C ADS SERVED</p> <p>0</p> <p>You are not bidding on any prompt</p>	<p>TRACKED COMPETITOR ADS</p> <p>0</p> <p>Datadog, New Relic, Splunk · all absent</p>	<p>"OTHER" ADVERTISERS</p> <p>2</p> <p>Monday.com, Sentry · not in your set</p>
---	---	--	--

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

ADVERTISER	DOMAIN	CARDS	PROMPTS HIT	CLASSIFICATION
Monday.com	monday.com	1	1	Adjacent: work / IT management pushing "AI agents handle incidents"
Sentry	sentry.io	1	1	Adjacent: error monitoring moving into AI root-cause and fixes

THE FIRST-MOVER WINDOW

No direct observability or SRE competitor is bidding on these prompts yet. A modest ChatGPT ad budget on your category 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 Datadog or New Relic notice the lane is open.

WHY IT STILL MATTERS THAT IT'S NEARLY EMPTY

Two adjacent vendors are already buying on your exact problem prompts. The same dynamic in adjacent B2B 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.

09B · THE ADS · VERBATIM

Read the two cards buyers already see **before** they see your name.

The sponsored ChatGPT carousel cards captured against your tracked prompts in today's audit, both served on non-branded SRE prompts where Vendor C is absent.

Monday.com

monday.com

Run IT better with AI agents

AI agents triage tickets, manage your KB, and handle incidents 24/7.

Served on: "how to build reliability memory so teams stop solving the same incidents twice"

Sentry

sentry.io

Debugging agent with prod context

Sentry's Seer finds the root cause, suggests a fix, and gets you back to shipping. Fast.

Served on: "best way to catch configuration drift before deployment"

THE COMPETITIVE READ

Neither **Monday.com** nor **Sentry** is an AI SRE platform. But both are paying to appear on the exact problem prompts Vendor C was built for, "reliability memory so teams stop solving the same incidents twice" is your Memory pillar, "configuration drift before deployment" is your Context pillar, and on both, Vendor C has zero organic presence. Adjacent tools are intercepting your buyer with your own message.

THE ASYMMETRY IN YOUR FAVOR

None of your tracked observability competitors are bidding, and none of the current advertisers are purpose-built for production control the way you are. A small, surgical ad spend on the "reliability memory" and "config drift" prompts would let you claim the paid layer on the exact queries that describe your product, while it's open.

CAVEAT WE'LL STATE PLAINLY

This is a same-day snapshot, so the ad layer is thin by definition. Treat two cards 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 fragile #1 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 3% non-branded gap.

- 1 Fix the brand disambiguation first**

ChatGPT asks "which Vendor C do you mean?" Resolve your entity: consistent Vendor C naming, Organization schema with sameAs links, a clear about / company page, and a Wikidata / Crunchbase footprint. Until the engines know who you are, nothing else compounds.
- 2 Win the non-branded SRE prompts you currently miss**

"Automated root cause analysis," "reduce MTTR with AI," "blast radius analysis," and "catch config drift" return zero Vendor C. Publish citable assets in that exact language, mapped to Context, Memory, and Control. This is the number that moves your real visibility from 3%.
- 3 Get named by the SRE and AIOps gatekeepers**

Reddit SRE and DevOps threads, CloudNativeNow, G2 AIOps category, and incident.io content are cited on your prompts and never mention you. Contribute, get into the round-ups, and earn citations alongside them.
- 4 Turn customers into proof the engines can quote**

Existing customer case studies (days-to-minutes outcomes) are already cited. Expand into named, metric-rich case studies and G2 reviews, so the engines have concrete outcomes to attribute to Vendor C, not just a homepage.
- 5 Own the "AI SRE / production control" language**

Google already describes you as "automated root-cause analysis built for engineering teams." Make Context, Memory, and Control the defining vocabulary of the category with definitive explainers, so the engines reach for your framing on generic prompts.
- 6 Take the open ChatGPT ad inventory**

Monday.com and Sentry are bidding on your exact problem prompts; no observability competitor is. Run modest sponsored cards on the reliability-memory and config-drift prompts to claim the paid layer while it's uncontested.
- 7 Run the audit monthly, and track both metrics**

Watch overall AND non-branded visibility every month, plus the Datadog mention gap and the disambiguation rate. The branded number flatters you; the non-branded number tells you whether you're actually winning the category.

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 C's founders pull this intelligence directly. No dashboards. No exports. Just ask.

STEP 1 · NOW

Workshop the playbook

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

STEP 2 · 30 DAYS

Expand the tracked set

Scale across root cause, MTTR, blast radius, alert fatigue, config drift, and runtime control prompts. Keep Datadog, New Relic, Splunk, Honeycomb, and Arize in the competitor graph.

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