ILLUSTRATIVE SAMPLE · SYNTHETIC DATA · BRAND AND COMPETITORS ARE FICTIONAL

AI SEARCH VISIBILITY DIAGNOSTIC

RelayHQ

AI customer support software · United States, SMB and startup · 29 October 2026

50%	10%	5	21
of questions name RelayHQ at least once	of questions put RelayHQ forward	questions where competitors appear and you do not	of 47 citations come from G2 and Capterra

Figures above are question-level — a question counts if the brand appears in any one of its three runs. Answer-level counts are on page 2.

What the answers say

Half the questions name RelayHQ; one in ten suggests it. The answers are assembled from review sites — 21 of 47 citations, against 26 from every other source type combined. **That points at category definition. It does not prove that changing the category would change the answer**, and the category is G2's decision rather than ours.

THREE FINDINGS · WHAT WE SAW, AND WHAT WE THINK IT MEANS

01 **In the consideration set, not in the answer**

OBSERVED — 5 of 10 questions name RelayHQ. In 4 of them it is one line in a five-tool list; in 1 it is put forward first.

OUR READING — The gap is between being listed and being chosen. Thirty answers show the gap; they do not show what produces it. Category placement is one candidate, not the only one.

02 **Review sites supply 21 of 47 citations**

OBSERVED — G2 is cited in 13 answers, Capterra in 8. Every other source type combined — editorial, community, vendor docs, Product Hunt — accounts for 26.

OUR READING — If review-site pages decide the shortlist, category placement inside them matters. But G2 sets its own taxonomy and eligibility rules, so this is a hypothesis to check before it is a recommendation.

03 **The price and Shopify questions have no RelayHQ in them**

OBSERVED — “Under $200 a month” and “best helpdesk for Shopify” are answered from review text and a competitor's app listing. Neither question names the product.

OUR READING — Both are channel gaps rather than content gaps: one needs a public price, the other needs a listing that does not exist. Neither is an AI-visibility project.

01 What the answers say 1
02 Question by question 2
03 What the answers actually said 3
04 What the answers cited 4
05 Where you are absent 5
06 Three things to do next 6
07 The plan, and when to stop 7
08 Method and limits 8

About this document. This is a design sample for a service proposal. The brand, its competitors and every number in it are fictional and synthetic, produced to show the structure of a real diagnostic — it is not a measurement of any real company. A live diagnostic replaces every value with the results of 30 logged-out AI answers, delivered with the answer-level evidence file described on page 4.

Question by question

10 QUESTIONS · 3 RUNS EACH · CHATGPT

Each question was asked 3 times in a logged-out session. A question counts as a mention if the brand appears in any one of the three answers, and as a recommendation if it is put forward as the suggestion rather than listed alongside others.

#	QUESTION PUT TO CHATGPT	RELAYHQ	BEACON	TRIAGE	LOOP
1	best AI customer support software for small teams	○	●	○	–
2	zendesk alternatives for startups	–	○	●	○
3	AI chatbot for ecommerce support under $200 a month	–	●	–	○
4	helpdesk with AI suggested replies	○	○	○	–
5	intercom alternatives 2026	–	○	●	○
6	best support tool with slack integration	●	○	–	–
7	AI support agent trained on your help center	○	○	○	–
8	cheapest helpdesk for a saas startup	–	–	○	●
9	customer support software with a good api	○	○	–	–
10	best helpdesk for shopify stores	–	–	–	○

● recommended ○ mentioned only – absent

Who the answers name

The pale bar is how often a brand is named at all. The solid bar is how often it is the one put forward. A brand can be everywhere and still never be the suggestion — that is the position RelayHQ is in.

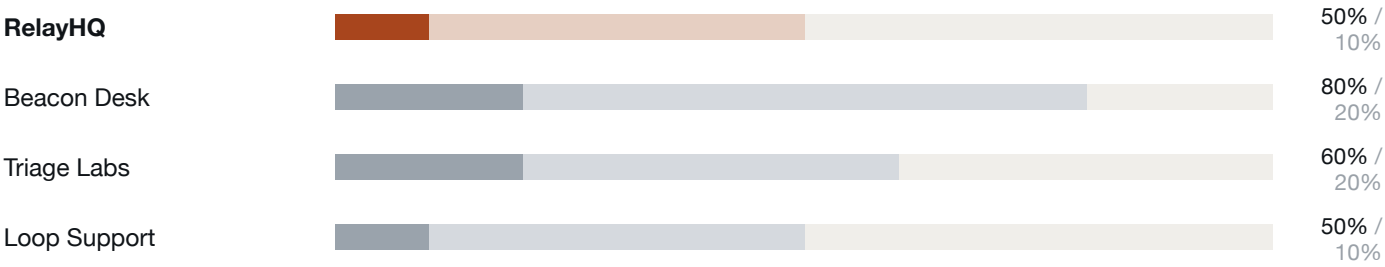

How to read the table

The right-hand half is the market. Where a competitor has a dot and you have a dash, that is a question the buyer asked and your brand was not part of the answer — 5 of them in this sample. Those are the content brief.

What the answers actually said

EXCERPTS · CHATGPT · 29 OCTOBER 2026

Four of the 10 answers, quoted. This is the part clients forward to their team, because it is the only part that sounds like a customer talking.

“best AI customer support software for small teams”

> For small teams the usual shortlist is Beacon Desk, Triage Labs and RelayHQ. Beacon Desk has the strongest review presence and the most integrations; Triage Labs is better value under 1,000 conversations a month; RelayHQ is often mentioned for Slack-first workflows. Most teams end up choosing on price at their actual volume.

What it tells us. RelayHQ is in the shortlist but placed third and framed as a niche. The framing comes from a G2 category page, not from anything the company wrote about itself.

“intercom alternatives 2026”

> Common Intercom alternatives include Beacon Desk, Triage Labs and Loop Support. Beacon Desk is the closest feature match; Triage Labs is cheaper for teams under five agents; Loop Support is the usual pick for Shopify-first stores. Pricing changes often, so confirm on the vendor site.

What it tells us. A named-competitor question answered entirely from two listicles. RelayHQ is absent from both, so it cannot appear here regardless of product quality — this is a distribution gap, not a positioning one.

“cheapest helpdesk for a saas startup”

> The cheapest options for a SaaS startup are usually Triage Labs (free up to 200 conversations) and Loop Support ($19 per agent). Free tiers matter more than headline price at this stage. Reddit threads on r/SaaS tend to recommend starting free and migrating once support volume stabilises.

What it tells us. Price questions are answered from free-tier limits quoted on Reddit. RelayHQ publishes neither a free tier nor a per-conversation price, so the model has nothing to compare.

“customer support software with a good api”

> If your team needs to build on top of the helpdesk, look at Beacon Desk and RelayHQ. Beacon Desk has the more complete REST API and webhooks; RelayHQ's API is smaller but simpler to integrate. Triage Labs has an API but limited documentation. Check rate limits before committing — they are usually where integrations break.

What it tells us. RelayHQ is named second and framed as “simpler”. That framing comes from a docs page that is indexed but thin — length and specificity are what the model is summarising.

> **The pattern to watch for**
>
> Every answer above is assembled from two or three sources, and the brand is missing from those sources — not from the answer. That is a lead about where to work, not a finding about what would change the answer.

What the answers cited

47 SOURCE APPEARANCES · 30 ANSWERS

The citations show which sources were visibly referenced in these answers; they do not establish that every claim was supported by them. Ordered by how many answers cited each source. It is the most actionable page in the report: these are the pages to be on.

SOURCE	TYPE	ANSWERS	WHAT IT CONTRIBUTES
g2.com	Review site	13	Category and comparison pages
review blogs	Editorial	9	“Best of” listicles, mostly affiliate
capterra.com	Review site	8	Category listings
reddit.com / r/SaaS	Community	7	Three pricing threads
vendor documentation	First-party	6	API and pricing pages
producthunt.com	Directory	4	Launch pages

Counted per answer: a source cited in three answers appears three times. 21 of 47 source appearances are review sites; the remaining 26 are spread across editorial, community, vendor documentation and directory pages. This is not a de-duplicated URL count.

Where the citations come from

Directories & review sites 53% · 25 Community 15% · 7 Editorial & blogs 19% · 9 First-party sources 13% · 6

Each cited source falls in exactly one category, and the counts sum to 47 — the same total as the table above. “First-party” means pages published by any brand in this category, not only the client's own.

53% of citations point at directories and review sites. That is the lever: being inside a list is a week of work; outranking the list is a year.

WHAT THIS LIST MAKES LESS LIKELY

Based on this sample, another blog post is not the clearest first move: no client-owned page is cited, and models rarely cite pages nothing else references. Whether a new page would earn citations remains uncertain.

WHAT IT POINTS AT

Two or three pages carry most of the citations, and none of them describes the brand. The plan on page 6 is about getting into those pages, not about outranking them.

> **Every row here has a paper trail**
>
> The 30 answers are delivered as a workbook: date, run number, the question, the answer text, each URL it cited, and how it was classified. If a client disagrees with a count on this page, we can point at the rows rather than argue.

SHEET 1
Prompts
The fixed 10, with the intent behind each one.

SHEET 2
Answers
One row per answer: prompt, run, date, text, sources, classification.

SHEET 3
Counts
Formulas that reproduce every percentage in this report.

Where competitors appear and you do not

5 GAPS · P1 = DO FIRST

These are questions a buyer asked where the answer named someone else. Each one is a page that does not exist yet, or a source that does not mention you.

01 **zendesk alternatives for startups** P1
In the answer: Beacon, Triage, Loop · **Cited:** g2.com, reddit.com
The answer is built from a G2 alternatives page. RelayHQ is not on it.
What to publish: Dual G2 category placement, so the alternatives pages can include you.

02 **AI chatbot for ecommerce support under $200 a month** P1
In the answer: Beacon, Loop · **Cited:** review blog, vendor pricing page
The price ceiling is the question. No per-conversation price is published, so the model cannot place the product under $200.
What to publish: A public pricing page with a per-conversation number the model can compare.

03 **intercom alternatives 2026** P2
In the answer: Beacon, Triage, Loop · **Cited:** g2.com, review blog
Named-competitor questions are answered from listicles. RelayHQ is absent from the top three.
What to publish: A presence in the two listicles that answer this question — outreach, not content.

04 **cheapest helpdesk for a saas startup** P2
In the answer: Triage, Loop · **Cited:** reddit.com, capterra
Reddit pricing threads name two tools by free-tier limits. RelayHQ is not quoted.
What to publish: A free tier, or a documented reason there is not one, quoted somewhere Reddit can repeat.

05 **best helpdesk for shopify stores** P3
In the answer: Loop · **Cited:** Shopify App Store, review blog
There is no Shopify app listing, so the answer has nothing to rank.
What to publish: A Shopify app listing.

ADD AT DAY 30
Named-competitor questions
“alternatives to X” — where a shortlist is already being built.

ADD AT DAY 30
Price-threshold questions
“under $X” — answered from pricing pages, not reviews.

ADD AT DAY 30
Segment questions
“for a small team”, “near me” — where the brand is usually missing.

Why we rank gaps, not keywords

A gap is a question with evidence that the answer exists, cites something specific, and names a competitor. That is a different and much smaller list than a keyword export — and it is the only list that maps to work you can finish in a month.

Three things to do next

RANKED BY EVIDENCE, THEN COST

Each recommendation carries the strength of the evidence behind it, who has to do the work, what it costs, and what result would show the reading was wrong. We do not rank by “expected effect per hour” — nobody can estimate that honestly from 30 answers.

P1 Publish the three-way comparison the questions are already asking for

EVIDENCE: DIRECT OBSERVATION | OWNER: YOU, PRICES VERIFIED BY US | COST: MEDIUM — 6 HOURS, PLUS VERIFICATION

THE AI HYPOTHESIS

Three of the ten questions are comparison questions answered from listicles. The cheapest way into that set is to be one of the sources.

VALUE REGARDLESS

Answers questions 2, 3 and 5; reusable in sales and support, not only in AI search.

WHAT TO DO

- One page: RelayHQ vs Beacon Desk vs Triage Labs, with price at 1,000 conversations, setup time and what each is bad at.
- State where RelayHQ loses. A comparison that claims total victory is neither cited nor believed.
- Verify every competitor price on the day of publication and date the page.

Live by day 30 · measured at day 60: page cited · Not claimed: Being the source does not guarantee a citation. We expect it to be a precondition, and the day-60 re-run is what tests it.

P2 Request the AI Customer Support category on G2 — if eligibility allows

EVIDENCE: REASONABLE INFERENCE | OWNER: YOU, WITH OUR CHECK ON THE CATEGORY RULES FIRST | COST: MEDIUM — 4 TO 6 HOURS PLUS A REVIEW CYCLE

THE AI HYPOTHESIS

The product is grouped under “Help Desk”, not “AI Customer Support”, so the buying questions never surface it. Dual placement is allowed for some vendors, not all.

VALUE REGARDLESS

Potentially the largest lever: review sites are 21 of 47 citations.

WHAT TO DO

- Check G2's category criteria and whether dual placement is permitted for this product.
- If eligible, request the category and complete the pricing fields — review sites summarise from them.
- Ask five customers for reviews that describe their use case, without suggesting wording.

Request filed by day 30 · measured at day 90: category live, answers may still not move · Not claimed: G2 decides. If the category is refused, the recommendation is dead and the report should say so at the re-run.

P3 Ship a Shopify listing and a public API page — as product work, not an AI spend

EVIDENCE: CONDITIONAL GAP | OWNER: YOUR PRODUCT TEAM | COST: HIGH — PRODUCT AND ENGINEERING TIME, NOT OURS

THE AI HYPOTHESIS

Two gaps are channel gaps. Neither is fixed by writing content, and neither should be funded from an AI-visibility budget on its own.

VALUE REGARDLESS

Opens two questions; only worth it if the integration is on the product roadmap anyway.

WHAT TO DO

- Confirm the Shopify integration exists and someone can maintain the listing.
- Publish an /api page with endpoints, rate limits and a working example.
- Link both from the comparison page so the sources connect.

Decided by day 30 · not measured in this engagement · Not claimed: If the integration is not real, do not list it. A thin listing is worse than none, and we will not write one.

WHAT WE WOULD NOT RECOMMEND

- Publishing more blog posts. No cited source in this sample is a brand blog.
- “AI-optimising” the homepage. None of the gaps is a messaging problem.
- Buying links or mentions. The sample cannot show they help.

IF YOU ONLY DO ONE THING

Do P1. It is the only recommendation backed by direct observation, it costs a few hours, and it has value outside AI search. The other two are worth doing only after P1's day-60 result is in.

The plan, and when to stop

TWO CLOCKS: WORK, THEN MEASUREMENT

Work has to land before day 30 for the day-60 re-run to mean anything. Anything that lands after day 30 is marked as such, and is measured at day 90 or not at all — we would rather say “too early” than score an action that had not happened yet.

WHEN	WORK	OWNER	LIVE BY	STOP / CHANGE IF
Week 1	Check G2 eligibility; verify competitor prices	Us + client	Day 7	Eligibility unclear — park P2, do not promise it
Week 2–3	Publish the comparison page	Client, with our draft	Day 21	Any unverified price — do not publish
Week 3–6	G2 review cycle, if eligible	Client	Day 42, if approved	Category refused — stop P2

Measurement calendar

DAY 30
Execution check
page live, eligibility answered, request filed
Out: Written confirmation, action by action

DAY 60
Effect re-run
the same 10 questions, 3 runs each
Out: Before/after table, same layout

DAY 90
Decision
continue, change the hypothesis, or stop
Out: Continue, change the hypothesis, or stop — in writing

> **The decision standard at day 90**
> At day 90 we compare each number against the range the three runs already produced. A change inside that range is noise. A G2 category approved in week 5 cannot show an effect by day 60 — for that action the honest answer is “too early”, not “failed”.

One row, end to end

The same trace exists for every row in the workbook. This is what a client sees when they ask “where did that recommendation come from?”.

QUESTION	P02 — “zendesk alternatives for startups”
ANSWER, RUN 2 OF 3	“Common Zendesk alternatives for small teams include Beacon Desk, Triage Labs and Loop Support. Beacon Desk is the closest feature match; Triage Labs is cheaper under five agents; Loop Support is the usual pick for Shopify-first stores.”
CITED	g2.com — “Zendesk alternatives” page · reddit.com/r/SaaS — pricing thread
CLASSIFIED	Mentioned: no · Recommended: no · Source appearances: 2
LEADS TO	**Gap 01 → P1. The answer is assembled from a list the brand is not on. The fix is to be one of the sources, not to be better than them.**

WHAT WE NEED FROM YOU

Two hours in week one, one person who can approve content, access to the analytics you already have, and an honest answer at day 30 about which actions actually shipped. Without the last one, the re-run measures the calendar, not the work.

WHAT ENDS THE ENGAGEMENT

If the day-90 comparison shows no change outside the baseline range, the honest recommendation is to stop paying for measurement and spend the budget on the underlying problem. We will say so in writing rather than propose another quarter.

Method and limits

SETUP

Model	**ChatGPT (consumer app, logged out)**
Prompts	**10 buying questions, fixed for the baseline**
Runs per prompt	**3, spread across 3 days**
Answers analysed	**30**
Location	**United States**
Dates	**26–29 October 2026**
Classification	**Mention / recommendation / citation, applied by hand**
Source counting	**Appearances per answer, not unique URLs**
Evidence file	**Answer-level log delivered with the report**

CLASSIFICATION

Mentioned — the brand name appears in the answer, in any position, including a list of ten.
Recommended — the answer puts the brand forward as the suggestion, or names it first.
Cited — the answer links or names a source; counted once per answer, not per link.

LIMITS YOU SHOULD KNOW BEFORE BUYING

- Answers vary between runs. Three runs show a range, not a fixed position; a single run is not evidence.
- Results are for logged-out, US, non-personalised answers. A logged-in buyer with a different history may see different tools.
- One engine only. Perplexity and Google AI Overviews draw on different sources and are quoted separately.
- G2 and Capterra rankings change monthly; the category audit is a snapshot of 29 October 2026.
- This is a 10-question sample of the category, not a census of everything ChatGPT says about RelayHQ.
- Correlation in 30 answers is a lead for the next test, not proof of what would change the answer.

WHAT THE $850 BUYS

The analysis: ten questions, thirty answers, the source audit, the ranked gaps, three evidence-ranked recommendations, the evidence workbook, a 60-minute walkthrough, a day-30 execution check and a day-60 re-run. The day-90 decision point is yours to take; we will prepare it, but nothing after day 60 is included.

WHAT IT DOES NOT BUY

Implementation. Publishing the pages, aligning the profiles, the review cycle, the G2 request, the Shopify listing — those are your team's hours or a separate quote. The report says which is which.

> **Sample notice.** Brand, competitors and all figures in pages 1–7 are synthetic and illustrate the format only. The method described on this page is the method used for a live diagnostic; the numbers on the previous pages are not a measurement of any real company.