

AI Search Visibility Audit: a well-known help desk tool

Category: help desk software · Anonymized public version: “Brand X” is the audited tool, “Competitor A” and “Competitor B” are two leading rivals · Data collected October 14, 2026 · 5 AI answer engines, 106 answers

Executive summary

When buyers ask AI assistants which help desk software to use, **Brand X is almost always on the list, but it is rarely the pick.** Engines name Brand X in 95 of 106 answers, more than Competitor A or Competitor B, and they describe it warmly. But they file it as “the simple, email-first inbox for small teams.” That label wins Brand X a mention and loses it the recommendation. On generic buyer questions, Competitor B is recommended 49% of the time and Brand X 34%.

35%

Brand X share of voice (all prompts), the highest of the three: Competitor A 31%, Competitor B 33%

34%

of generic category answers recommend Brand X vs 49% Competitor B and 38% Competitor A

74%

of answers naming Brand X describe it as an email-first or inbox tool (70 of 95) vs 8% for Competitor A

0/12

Perplexity category answers that recommend Brand X (Competitor A 6/12)

Three things that matter most


- **Known everywhere, chosen less often.** Brand X is named in 95 of 106 answers and has the best tone of the three (net sentiment +0.71 vs +0.27 Competitor A and +0.57 Competitor B). But when an engine picks one tool, it is usually Competitor B for small businesses (“my default pick”) and Competitor A for complex teams. Brand X gets the “if you mostly handle email” slot.
- **It drops out of the “leaders” questions.** In the “what do experts recommend” and “list the leading vendors” questions, Brand X appears in 12 of 18 answers, while Competitor A appears in 18 and Competitor B in 17. Engines treat it as a niche choice, not a category leader. 19 answers attach a limit to it (“lightweight”, “basic”, “lacks SLA features”), compared with 8 for Competitor A.
- **Other vendors write the lists engines read.** 7 of the 10 most-cited pages are “best help desk software” guides or product pages published by help desk vendors. The single most-cited page is another vendor’s own “best help desk software” guide (25 citations). Brand X’s own pages are cited 53 times, but 21 of those are pricing and plan pages.

Share of voice vs competitors, by engine


Share of voice is each brand's share of all mentions of the three tracked brands. The first chart covers all 12 questions. The second covers only the 7 generic category questions, where no brand is named in the prompt.

Brand X Competitor A Competitor B

All questions: share of voice (%)



Generic category questions only: share of voice (%)



Engine	Brand X SoV	Competitor A SoV	Competitor B SoV	Brand X named, category Qs	Brand X recommended
ChatGPT (OpenAI)	35%	32%	33%	13/14	13/24
Claude (Anthropic)	36%	30%	34%	13/14	4/24
Gemini (Google)	35%	32%	32%	12/14	14/24
Perplexity	33%	31%	37%	7/12	2/22
Google AI Overviews	38%	35%	27%	6/7	6/12

Question by question (answers naming each brand)

Question	Type	Brand X	Competitor A	Competitor B
What are the best help desk software tools?	category	9/9	9/9	9/9
Which help desk software would you recommend for a small business?	category	6/8	5/8	8/8
What is the best help desk software for startups in 2026?	category	8/8	7/8	8/8
Top alternatives to Competitor A for help desk software?	alternatives	9/9	9/9	9/9
Top alternatives to Competitor B for help desk software?	alternatives	9/9	9/9	9/9
Brand X vs Competitor A: which is better?	branded	9/9	9/9	0/9
Brand X vs Competitor B: which is better?	branded	8/9	1/9	8/9
Is Brand X a good help desk software?	branded	9/9	8/9	6/9
What help desk software do experts recommend?	category	6/9	9/9	8/9
Cheapest reliable help desk software?	category	7/9	2/9	8/9
Which help desk software has the best customer support?	category	9/9	8/9	8/9
List the leading help desk software vendors and their strengths.	category	6/9	9/9	9/9

Each question was asked twice per engine on ChatGPT, Claude, Gemini and Perplexity, plus one Google AI Overview check, for up to 9 answers per question. Red marks a question where the brand never appeared. Branded questions (“Brand X vs ...”, “Is Brand X good...”) put Brand X in the prompt, so its numbers there are inflated by design.

Recommend rate and sentiment

Recommend rate (all 106 answers)



Mention rate (answers that name the brand)



Recommend rate on generic category questions



Mention rate on generic category questions



An answer counts as recommending a brand if it ranks the brand first in a list or uses recommendation language (“best”, “I’d recommend”, “go with”) with positive wording about it. On the 61 generic answers, Brand X is recommended 21 times, Competitor A 23 and Competitor B 30. Perplexity never recommended Brand X on a generic question, and Claude did so twice.

Brand	Answers naming it	Positive	Neutral	Negative	Net sentiment	Email-first framing
Brand X	95	69	24	2	+0.71	70 (74%)
Competitor A	85	33	42	10	+0.27	7 (8%)
Competitor B	90	54	33	3	+0.57	13 (14%)

The tone toward Brand X is the warmest of the three. The issue is the box it is put in. Typical wording:

"Pick [Brand X] if you have a small support team and mostly handle email."

ChatGPT (OpenAI) · "What are the best help desk software tools?"

"[Brand X] : Best for email-driven support teams wanting a shared inbox experience."

Google AI Overviews · "What help desk software do experts recommend?"

"Not Ideal for Complex Workflows: If you need deeply intricate routing rules, multi-layered ticket hierarchies, or massive enterprise customizations, [Brand X] may feel a bit too lightweight."

Gemini (Google) · "Is [Brand X] a good help desk software?"

"[Brand X]'s lack of SLA features could be a dealbreaker for larger support teams."

Claude (Anthropic) · "[Brand X] vs [Competitor A]: which is better?"

Sentiment is scored automatically with a word-list method; phrases like "lightweight" or "lacks SLA features" are not counted as negative, so the limits are counted separately above.

The sources AI cites, and who writes them

All five engines cited web sources in this audit. Brand X is missing from only 11 of 106 answers, so the question here is less "which sources leave Brand X out" and more "which sources decide how it is framed." The table shows how often each domain was cited in answers without and with Brand X.

Cited domain	Type	Total answers citing	Brand X absent	Brand X present
<i>Other help desk vendor #1 site</i>	Other help desk vendor	31	6	25
<i>Brand X site</i>	Brand X-owned	31	0	31
<i>Other help desk vendor #2 site</i>	Other help desk vendor	30	2	28
<i>Other help desk vendor #3 site</i>	Other help desk vendor	30	5	25
<i>Other help desk vendor #4 site</i>	Other help desk vendor	25	7	18
g2.com	Review site	24	2	22
<i>Competitor A site</i>	Tracked competitor	23	4	19
<i>Other help desk vendor #5 site</i>	Other help desk vendor	18	1	17
<i>Competitor B site</i>	Tracked competitor	17	3	14
thecxlead.com	Media / review publisher	17	1	16
<i>Other help desk vendor #7 site</i>	Other help desk vendor	13	0	13
<i>Other help desk vendor #6 site</i>	Other help desk vendor	13	2	11
<i>Other help desk vendor #8 site</i>	Other help desk vendor	13	2	11
<i>Other help desk vendor #9 site</i>	Other help desk vendor	12	0	12
<i>Other help desk vendor #10 site</i>	Other help desk vendor	10	0	10
<i>Other help desk vendor #11 site</i>	Other help desk vendor	10	1	9

What this shows

- **Vendor-written "best of" lists dominate.** Guides and blogs published by at least five other help desk vendors are cited more than most independent reviews. Each vendor puts itself first and describes the others in a line, which is where "simple email inbox" comes from. One Claude answer says so directly:

“Several 2026 roundups point to similar tools for each need, though many of these lists are published by vendors or affiliate sites, so weigh them accordingly.”

Claude (Anthropic) · “What help desk software do experts recommend?”

- **G2 is the independent source that matters most.** G2's help desk category page and its “best help desk software” article are cited 36 times. thecxlead.com, thectoclub.com and Software Advice follow.
- **Brand X's cited pages describe price, not position.** The pricing page alone is cited 16 times. Its comparison pages (vs Competitor A, Competitor B and another vendor) and “Competitor A alternatives” post are cited 52 times, which helps in head-to-heads. No Brand X page is cited for “best help desk software” or “leading vendors.”

Most-cited individual pages (all engines)

URL (tracking parameters removed)	Citations
[Other help desk vendor #3]/learning-center/best-help-desk-software	25
g2.com/categories/help-desk	23
[Other help desk vendor #1]/learn/top-help-desk-softwares/	20
[Other help desk vendor #4]/blog/best-help-desk-software/	16
[Brand X]/pricing/	16
learn.g2.com/best-help-desk-software	13
[Other help desk vendor #8]/blog/best-help-desk-solutions	9
[Competitor B]/helpdesk/software/	9
[Other help desk vendor #2]/blog/best-help-desk-for-startups	9
[Other help desk vendor #22]/blog/best-help-desk-software-top-picks-comparison	7

Gemini citations are recorded by source domain. YouTube video URLs are excluded from the page list.

Engine-by-engine notes

ChatGPT (OpenAI, gpt-5.6-luna, web search on)

ChatGPT names Brand X in 13 of 14 category answers and recommends it in 7. Competitor B gets 11. Its tables usually give Competitor B “best overall” and Brand X “best for simple, human-centered support.”

“For most small businesses, I'd recommend [Competitor B].”

ChatGPT (OpenAI) · “Which help desk software would you recommend for a small business?”

Claude (Anthropic, claude-haiku-5-5, web search on)

Claude's searches worked in this audit. It names Brand X in 13 of 14 category answers but recommends it in only 2, and it calls Competitor B “my default pick” for small businesses.

“[Competitor B] (my default pick): Several guides describe it as approachable for small teams.”

Claude (Anthropic) · “Which help desk software would you recommend for a small business?”

Gemini (Google, gemini-3.5-flash-lite, Google Search grounding)

Gemini is the most balanced engine: it recommends Brand X in 14 of 24 answers, close to Competitor A (13) and Competitor B (12). It also gives the fairest description of Brand X's channels beyond email.

"[Brand X]: Great for email-first teams, but also handles live chat ([chat widget]), basic social media integrations, and knowledge bases ([knowledge base product])."

Gemini (Google) · "[Brand X] vs [Competitor B]: which is better?"

Perplexity (perplexity/sonar)

Brand X's weakest engine. Perplexity names it in 7 of 12 category answers and never recommends it there. Its "quick picks" lists sometimes leave Brand X out entirely:

"Quick picks: Choose [Competitor B] for a practical starting point, [Competitor A] for complex or high-volume operations, [another vendor] for in-app messaging and AI-led support, or [another vendor] if cost is a priority."

Perplexity · "What are the best help desk software tools?"

Google AI Overviews (via Google Search, US English)

An AI Overview appeared for 11 of 12 queries. AI Overviews are Brand X's best channel: it is named in 6 of 7 category overviews and recommended in 5, usually as "best for small businesses."

"[Brand X]: Best for small businesses and human-centered support, featuring a clean, email-style shared inbox."

Google AI Overviews · "What are the best help desk software tools?"

Prioritized fix list

P1 Earn a “leader” framing, not just a “simple inbox” one. Brand X is missing from a third of the “experts recommend” and “leading vendors” answers. Publish crawlable, plain-text proof of scale and breadth (customer count, agent count, channels, AI features, uptime, review scores), and pitch it to the independent roundups engines already cite (G2’s category and article, thecxlead.com, thectoclub.com, Software Advice).

P1 Answer the “best help desk software” question yourself. Other help desk vendors own the most-cited pages, and they frame Brand X in one line. A fair, well-sourced guide on its own site (“best help desk software for small and growing teams”, “...for SaaS”, “...for ecommerce”) gives engines a Brand X-authored source for the questions where it is currently described by rivals.

P1 Close the gap on the “limits” engines repeat. Engines say Brand X “lacks SLA features”, is “lightweight” for complex workflows and has a “limited” free plan. Where these are outdated, publish dedicated feature pages (SLAs, routing, workflows, AI agent, reporting) with concrete specifics, and update the comparison pages engines already cite.

P2 Give Perplexity and Claude something to recommend. Both name Brand X but rarely pick it. They lean on short “quick picks” summaries from roundups. Getting a clear “best for [use case]” line on high-authority lists, and stating that use case consistently on the homepage and pricing page, makes the recommendation easier to repeat.

P2 Keep the pricing page doing double duty. It is cited 16 times. Add a short, plain-text “who each plan is for” section and keep plan names identical across docs and pricing, so engines quote one consistent price and position.

P3 Technical baseline. Make sure AI crawlers (GPTBot, OAI-SearchBot, ClaudeBot, PerplexityBot, Google-Extended) can reach pricing, compare and feature pages. Add Organization, SoftwareApplication/Offer and FAQPage structured data.

Methodology and limits

- **Date and scope:** collected October 14, 2026 (ET), US English. 12 buyer questions in three groups: 7 generic category, 2 “alternatives to [competitor]” and 3 Brand X-branded. Tracked brands: the audited tool (Brand X) and two of its best-known direct competitors (Competitor A, Competitor B). Names, domains and identifying details are withheld in this public version.
- **Engines and models:** OpenAI gpt-5.6-luna with web search (max 1 search per answer); Anthropic claude-haiku-5-5 with web search (max 1); Google gemini-3.5-flash-lite with Google Search grounding; Perplexity perplexity/sonar with web search; Google AI Overviews collected from Google Search results via DataForSEO.
- **Sample size:** 2 samples per question per engine for the four chat engines, plus 1 Google AI Overview per question: 106 answers in total (2 Perplexity requests failed with rate-limit errors and were not retried; 1 query returned no AI Overview). This is a small sample. AI answers vary from run to run, so differences of a few points aren’t meaningful; the large gaps (recommend rate, email framing, the experts/leaders questions) are robust in direction.
- **Analysis:** mentions, list position, recommendations, citations and sentiment were extracted with deterministic rules applied to the saved answers. No AI model was used to score or summarize. Quotes are verbatim (formatting removed). Share of voice counts only the three tracked brands; several other vendors appear often but were not

tracked. “Email-first framing” counts answers where a sentence naming the brand also says email-first/centric/based/driven, shared inbox or inbox.

- **Known limits:** word-list sentiment misses phrases like “lightweight”. List-position metrics are unreliable for answers formatted as tables, and the “alternatives to Competitor A/Competitor B” questions name the competitor in the prompt. API answers can differ from the consumer apps (which add personalization, location and history).

Prepared by Answerweight as a sample audit using publicly available AI answers. The audited company and its competitors were not contacted and did not take part. In this public version, brand names, domains, product names and other vendor names have been replaced; quotes are otherwise verbatim, and replaced words are shown in [brackets].