

The buyer who asked an assistant and never found you

Nobody called, nobody visited, and nobody said no. A buyer asked a question somewhere you cannot see, and your firm was not in the answer.

Author	Amit Govil, Founder and CEO, XONIK
Published	October 22, 2026
Reading time	8 min read
Online	https://xonik.com/insights/perspectives/the-buyer-who-asked-an-assistant-and-never-found-you/

IN BRIEF

A buyer who asks an AI assistant for suppliers and does not see your name leaves no trace: no visit, no form, no note. Every record a business keeps starts with a contact, so the loss never reaches them. Only a record you build yourself, of what assistants answer to your buyers' questions, shows whether you were left out.

KEY FINDINGS

1. A shortlist is the end of a chain the buyer never sees.

The assistant chooses which searches to run, which pages to open and which passages to keep, and the vendors describe each of those steps in their own documents.

2. Every record a business keeps begins with a contact, so a buyer who made none leaves nothing.

Your reports show the quiet a slow market would show, and from inside the data that quiet cannot be told from a market that is simply slow.

3. Server logs and robots.txt cannot tell you what was read on a buyer's behalf.

A visit by a user agent shows that a page was fetched, not the buyer's question or identity. Perplexity says its user fetcher generally ignores robots.txt, and OpenAI says robots.txt may not apply to its own.

4. Search Console shows where you were present, not where you were absent.

Google reports appearances in its AI features within the Web search type. That says nothing about the buyer in whose answer you were missing, or about assistants that are not Google's.

5. A question log, a question at first contact and a read of your logs and firewall are the means you have.

None secures a place in any answer. They remove the ways of being absent by accident and let you see the result, and a failed test may reflect settings on the buyer's side or the model's recall rather than a

failed page.

How does a shortlist form inside an answer?

A buyer who once typed a short phrase into a search box now writes a paragraph. It carries the situation, the constraint and the comparison: who could rebuild how our group bookings are handled, what the options are, and which suits a business of our shape. The assistant does not hand back a page of links to sort. It returns a composed answer with names in it, and the names are the shortlist.

The vendors describe how those names are reached. Google says its AI Overviews and AI Mode may use a technique it calls query fan-out, which issues multiple related searches across subtopics and data sources to build one response, and that AI Mode is meant for questions that need exploration, reasoning or comparison (Google Search Central, "AI features and your website", last updated 10 December 2025). Anthropic's documentation for its web search tool says Claude searches when a request concerns specific organizations, people or products that might have changed, and answers directly, without searching, when the request draws on stable knowledge. OpenAI's web search guide says the model may choose whether to search from the content of the prompt, and lists separate actions for searching, opening a page and finding text inside a page.

So a shortlist is the end of a chain the buyer never sees: the searches the assistant chose to run, the pages it opened, the passages it kept. A supplier missing from that chain was not rejected in any sense a person would recognize. It was never in the pile.

A supplier missing from that chain was not rejected in any sense a person would recognize. It was never in the pile.

What can you see of that buyer?

Almost nothing, and the reason is structural. Every record a business keeps begins with a contact. Analytics begins when a visitor arrives, a form when someone fills it, a sales note when someone speaks. A buyer who read an answer, formed a shortlist without your name and called other firms creates none of these. Your reports show the quiet they would show in a slow market, and from inside the data the two cannot be told apart.

Exhibit 1 follows one question and marks where your own record ends. The only place a trace can form is on the assistant's side of the line. OpenAI documents ChatGPT-User, Anthropic documents Claude-User and Perplexity documents Perplexity-User as agents that may visit a page when a person asks a question, and each publishes how its agent identifies itself. A visit by one of them can therefore appear in your server logs. It shows that a page was fetched. The pages we read describe the agent's name and its published address ranges; none says that the request carries the buyer's question or the buyer's identity.

EXHIBIT 1

Where a buyer's question leaves your record



Everything before the red marker happens outside your systems. The first thing your records could show is a contact, and for this buyer none comes.

Why is a block in robots.txt not a way to know?

Two further limits apply. Perplexity states that its user fetcher generally ignores robots.txt rules because a person asked for the fetch, and OpenAI says robots.txt rules may not apply to ChatGPT-User for the same reason. A rule in that file is therefore not a dependable way to learn, or to control, what is read on a buyer's behalf. And the answer the buyer saw is stored nowhere you can reach. Google says AI Mode and AI Overviews may use different models and techniques, so the responses and links they show vary. The question you test today is a sample, not a replay.

Which of your records can tell you anything?

Exhibit 2 sets out what the records a business usually keeps can and cannot show about a buyer who was never in contact. Google says appearances in its AI features are included in the overall search traffic reported in the Performance report, under the Web search type. That tells you where you were present. It cannot tell you about the buyer in whose answer you were absent, and it says nothing about assistants that are not Google's.

EXHIBIT 2

What each record you keep can and cannot tell you about a lost buyer

Record	What it holds	What it cannot show
Web analytics	Visits that arrived, and what they did	A buyer who never arrived
Search Console	Google's Performance report, which Google says includes appearances in its AI features	Any other assistant, and any buyer in whose answer you were absent
Server logs	Requests from fetchers such as ChatGPT-User, Claude-User and Perplexity-User	The question asked and the person who asked it

Sales notes	Reasons given by buyers who spoke to you	Buyers who never did
A question log you keep	What assistants answer to your buyers' questions, on the days you ask	What a particular buyer saw on a particular day

The last row is the only record that starts before a contact, and a business has to build it for itself.

What does a lost buyer look like from inside the business?

Take an operations director at a hospitality group who needs a partner to rebuild how group enquiries are handled. She writes her situation to an assistant and gets an answer naming firms. Two sound right. She calls them. The firm that has done this work for operators like hers is not named, because the page that describes the work is written in the firm's internal vocabulary and never uses her words.

From inside that firm the quarter looks like drift. First calls come later than expected, or from a different mix of buyers, and the shortfall is put down to timing, to a competitor's campaign, to the market. In our experience the calls that do arrive carry a tell: the buyer uses phrases that came from an answer, and the firm's own wording is absent from the conversation. That is a sign the first conversation has moved somewhere the firm does not measure. It is a pattern we have seen in delivery work, not a measured rate.

Suppose that firm puts the director's question to an assistant and opens the pages the answer links to. Each begins with the buyer's problem, in the buyer's terms, and says plainly what the firm does about it. Its own page begins with a statement of outlook. Nothing was wrong with its reach or its eligibility. What it lacked was a page that matched the question the director asked, and the first time it learned that was by reading someone else's.

What else shapes the answer that you never see?

Some of what decides the answer sits on the buyer's side of the screen. Anthropic's documentation says that an administrator can restrict which domains its web search tool may search, and OpenAI's web search guide lists hosted controls such as filters and sources. A buyer working inside a company account may be using an assistant configured to look only where the company allows. You cannot inspect those settings, and no change to your pages reaches a buyer whose assistant was told not to look.

The same documents show a second case. A model is not obliged to search. When it does not, the names in the answer come from the model itself, and a page you rewrite today cannot change what it recalls. The vendor pages we read do not say what decides whether a firm is recalled. We mention the case so that a failed test is not read as a failed page.

What changes it?

Three habits, none of which needs a new tool. First, keep a question log. Write down the questions your best buyers asked before they first called, in their words. Put each to the assistants your buyers are likely to use, on a fixed day each month, and file the answer with its links, the names that appeared and how each was described. That file is the record no other system keeps.

Second, ask. At the first conversation, ask a new contact what they asked and which tool they used before they reached you, and record it beside the lead source. Third, read your logs and your firewall. Perplexity's

page says a web application firewall may need to allow its agents explicitly. Anthropic's says that blocking by address can stop its crawler from reading robots.txt, so an address block is not a reliable way to state a choice. Google names the condition for a link in its AI features: the page must be indexed and eligible to be shown with a snippet, with no additional technical requirement.

None of this secures a place in any answer. It removes the ways of being absent by accident, and it lets you see the result.

What does this rest on?

Vendor documents read on 20 October 2026, and our delivery experience of how buyers describe their search to a firm. We did not test any assistant for this piece, and it says nothing about the number of buyers who use assistants, how often a name is dropped, or what any assistant would say about your firm.

1. [Google Search Central, AI features and your website: query fan-out, AI Mode, and where appearances are reported](#)
2. [OpenAI, Overview of OpenAI crawlers: ChatGPT-User and how robots.txt applies to it](#)
3. [OpenAI, Web search guide: when the model chooses to search, the actions it can take and the hosted controls](#)
4. [Anthropic, Does Anthropic crawl data from the web: Claude-User and blocking by address](#)
5. [Anthropic, Web search tool: when Claude searches, and restricting domains](#)
6. [Perplexity, Perplexity crawlers: Perplexity-User, robots.txt and firewalls](#)

USE THIS

[Read the research on being found through assistants](#)

[Use the AI search visibility checklist](#)

Today, ask one assistant the question your best client asked before they first called you, and save the answer.

This document was made from the article at

<https://xonik.com/insights/perspectives/the-buyer-who-asked-an-assistant-and-never-found-you/>. The online article is the current version.