Wynter

62%

How B2B buyers research vendors with AI.

62% of B2B buyers open their AI vendor research with a plain category question that names no vendor, even though 80% already have one in mind. We asked 200 B2B buyers what they typed, what they asked next, and what they did after the chat.

200	44%	13	Oct 2026
B2B BUYERS	AT 1,000+ EMPLOYEE COMPANIES	QUESTIONS	FIELDED

THE SHORT VERSION

The AI builds the shortlist. Your website has to confirm it.

62%

of first prompts are a category question with no vendor named.

67%

say the AI put a vendor they hadn't considered on their shortlist.

75%

sent 3 or more follow-up messages after the first answer.

44%

of picks after an open prompt and follow-ups weren't in the AI's first answer.

72%

went to vendor websites after the AI chat.

01	They ask for the category
02	The AI adds vendors
03	Thin prompts
04	The buying happens in follow-ups
05	The first answer isn't final
06	The AI picks and drafts
07	Then they visit websites
08	The whole journey
09	What to do with this

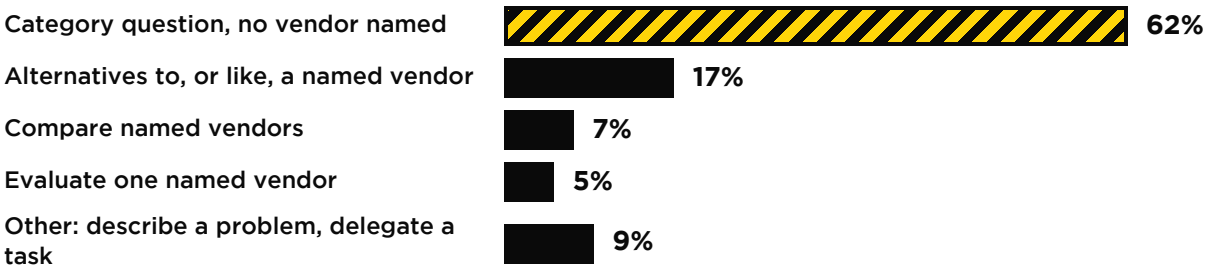
01 THE FIRST PROMPT

Buyers know vendor names but ask AI for the category.

62%

of first prompts are a category question

Four in five buyers (80%) walked into the chat with at least one vendor already in mind. Half had two or three. Then most of them typed a question that left those names out.



What the first message to the AI asked for, coded from the message each respondent wrote or pasted. Share of respondents (n=200).

Of the buyers who had a vendor in mind, **64% named none** in their first prompt. One engineer explained why:

> **"I started broad. To get a list without my influence."**
>
> SENIOR SOFTWARE ENGINEER, 50-199 EMPLOYEE COMPANY

The typical opener is a category, a use case and one requirement: "Recommend an intent data platform for our company to purchase. It should be easy to use, provide data on in market buyers, and help our sales team with outreach." When a vendor does appear, the most common pattern is "alternatives to X".

STAGE WHEN THE CHAT STARTED	SHARE OF BUYERS	NAMED A VENDOR IN THE FIRST PROMPT
Exploring the category	52%	28%
Building a shortlist	40%	33%
Making a final decision	8%	60%

Even buyers who were already building a shortlist mostly asked the category question. Only at the final decision do vendor names lead, and few buyers start an AI chat that late (n=15).

What it means: being remembered gets you into the buyer's head. It doesn't get you into the answer. The answer is built from how the AI understands the category, and whether it connects you to the problem the buyer describes in their own words.

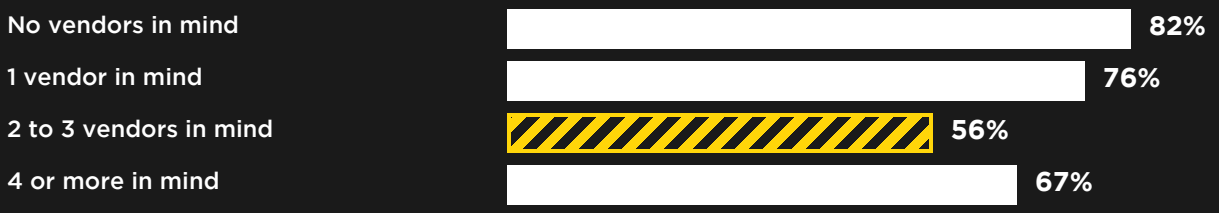
02 THE SHORTLIST

The AI adds a vendor to most shortlists.

67%

got a new vendor onto their shortlist from the AI

Because buyers ask for the category, the AI answers with the category as it sees it. Two in three buyers (67%) say it introduced a vendor they hadn't considered, and that vendor made their shortlist.



Share who say the AI introduced a new vendor that made their shortlist, by how many vendors they had in mind before the chat (n=198; 4 or more, n=15).

The surprising number is the middle one. More than half of buyers who already had two or three vendors in mind still left with a new name. Being on a buyer's list before the chat starts doesn't make the list final.

It also isn't only early-stage browsing. 75% of buyers who were actively building a shortlist got a new vendor from the AI, more than the 64% who were just exploring the category.

> *"Give me a list of new and upcoming vendors in the A/B testing space to consider. We are looking to evaluate new vendors, not existing vendors which have been in the space for a while."*
>
> LEAD WEB STRATEGIST, CRO, ENTERPRISE

> *"Then a comparison table with new candidates I never heard of."*
>
> STAFF SOFTWARE ENGINEER, 50-199 EMPLOYEE COMPANY

What it means: the consideration set is now co-written by the AI. Incumbents of mind can be displaced, and unknown vendors can be added, by how well the AI matches each vendor to the buyer's problem.

03 BUYER CONTEXT

First prompts are thin, and buyers think they said more.

1

buyer constraint in the median first prompt

We coded every first prompt for the buyer details it contained: company size, industry, budget, current tool, tech stack, region or compliance, role, and must-have features. The median prompt had **one**. 30% had none. Only 12% had three or more.

Then we asked the same buyers which of those details their first message included. They claimed 1.8 times more than they wrote. Relative to what they wrote, the gap is widest for who they are: role (5x) and industry (2.6x).

DETAIL	FOUND IN THE PROMPT THEY WROTE	CLAIMED ON A CHECKLIST
Must-have features	34%	63%
Tech stack or integrations	21%	39%
Industry	10%	26%
Their role	4%	22%
Budget	17%	20%
Current tool	12%	18%
Company or team size	14%	18%
Region or compliance	7%	12%

Written = coded from the first message each respondent wrote or pasted. Claimed = the same respondents' answer to "which of these did that first message include?" (n=200).

Some of that context lives in the tool instead. Between 30% and 46% started with a project, saved instructions, uploaded files or an existing chat. Their written prompts were no richer than everyone else's (median one constraint for both groups), which suggests the tool may have carried context they never typed.

4%

told the AI to take on a persona ("act as a cloud architect")

4%

asked for sources in the first prompt

28%

asked for a format: a table, a top 3, pros and cons

What it means: the AI fills the missing context with assumptions. Who a vendor is for, which company size it fits, what it replaces: if that isn't clear in what's published about you, the AI guesses. And the buyer may not notice the guess, because they believe they told it.

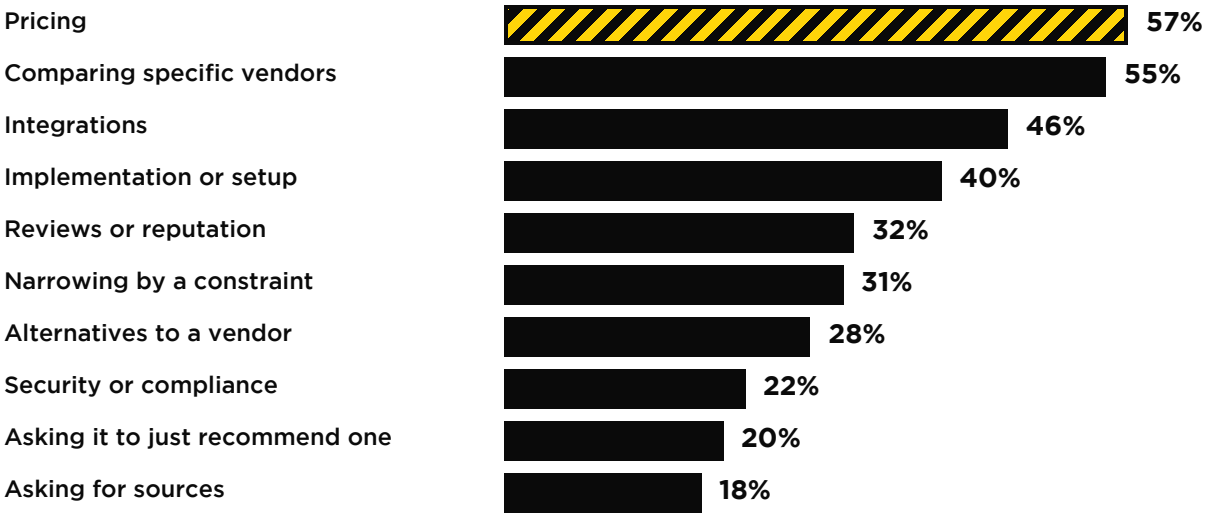
04 THE FOLLOW-UPS

The real buying happens in the follow-ups.

75%

sent 3 or more follow-up messages

Almost nobody stops at the first answer. 75% sent three or more follow-ups and 34% sent six or more. Only 4% sent none. The follow-ups are where the thin first prompt gets its real requirements.



What respondents' follow-ups asked about, multiple answers allowed. Share of respondents (n=200).

The messages themselves are short and blunt: a median of 13 words each. The first follow-up most often removes vendors. It adds a filter (19%) or checks price (17%). About half of conversations (52%) add a constraint or a fact about the buyer that the first prompt left out.

> **"We have a budget of $500/month. Please eliminate the vendors that don't fall within this limitation. Eliminate any vendors that don't have an API integration with Marketo."**
>
> ABM CAMPAIGN MANAGER, ENTERPRISE

> *"Please drop any companies not based in US. Please clean up list and stack rank based on reviews."*
>
> BUSINESS DEVELOPMENT REPRESENTATIVE, 200-999 EMPLOYEE COMPANY

Vendor names arrive here too. 31% named a vendor in the first prompt, and another 21% typed a vendor name only in a follow-up, often one the AI had just listed. Of the conversations that mention a vendor by name at all, four in ten do it only after the AI has answered.

WHAT THE FOLLOW-UP DID	FIRST FOLLOW-UP	LAST FOLLOW-UP
Filter or add a constraint	**19%**	8%
Ask about price	**17%**	11%
Compare or rank vendors	13%	12%
Setup and integration	4%	**13%**
Turn it into a table or doc	6%	**11%**
Ask for reviews or sources	4%	**10%**

Coded from the follow-up messages respondents wrote out, in order (n=114 conversations with two or more follow-ups). Conversations open by cutting the list and close with proof, setup and paperwork.

What it means: the decisive questions are the buyer's filters: budget, integrations, team size, region, compliance, the tool being replaced. A vendor survives the conversation when the AI can answer those questions about it clearly and correctly.

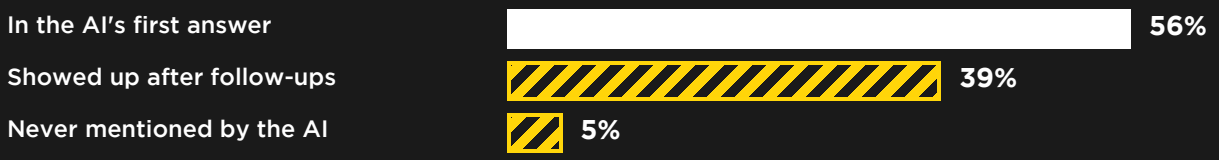
05 THE FIRST ANSWER

Four in ten winners weren't in the first answer.

44%

of picks after follow-ups surfaced later, or never

We asked buyers whether the vendor they chose was in the AI's first answer. Among those who started with an open prompt, followed up, and picked a vendor, **56%** said yes. **39%** said it only showed up after follow-ups, and **5%** said the AI never mentioned it.



Where the chosen vendor first appeared. Buyers who picked a vendor after an open first prompt and at least one follow-up (n=110). Excludes buyers whose first prompt asked to compare or evaluate specific named vendors, and buyers who sent no follow-ups.

The longer the conversation, the more likely the winner came from later in it. With two follow-ups or fewer, 29% of picks surfaced after the first answer. With six or more, it was 51%.

FOLLOW-UPS SENT	BUYERS	PICK SURFACED AFTER THE FIRST ANSWER, OR NEVER
0 to 2	21	29%
3 to 5	52	44%
6 or more	37	**51%**

> **"My colleague has activated a trial for this other option. Why was this not on your initial list?"**
>
> MARKETING MANAGER, 10-49 EMPLOYEE COMPANY

> *"Most of the companies you mentioned are brand level and don't let me track individual products."*
>
> SENIOR MACHINE LEARNING ENGINEER, 10-49 EMPLOYEE COMPANY

Treat these numbers as directional. 41% of the buyers who said they "chose" a vendor were still exploring the category, so "chose" often means "leaning toward". And being in the first answer can mean being one name in a list of ten.

What it means: the first answer is a starting line. Tracking a fixed set of one-shot prompts measures the opening of a conversation that usually runs three to six turns. The vendors that win are the ones that still fit after the buyer adds their real constraints.

06 THE DECISION

Buyers ask the AI to decide, and to write it up.

30%

asked the AI to pick a vendor, by their own report or their messages

One in five buyers (20%) say they asked the AI to just pick one. Counting what they wrote in their follow-ups too, it's three in ten (30%). Some explicitly asked it to choose using what it already knew about them or their company. No prompt tracker can see that context.

> *"Between these options do you think ClickUp sounds best based on what you know about my work?"*
>
> SENIOR CONTENT WRITER, ENTERPRISE

> *"If you have to pick one vendor today, what will you pick based on the company immediate needs for the sales team right now, given the facts."*
>
> SOFTWARE ENGINEER, ENTERPRISE

And the chat output often becomes the internal buying document. 36% asked for a specific format at some point: a table, a scorecard, a ranked top 3, pros and cons. About one in twenty asked the AI for something meant for colleagues or leadership: a decision memo, an executive summary, a board pitch, a shareable spreadsheet, slides.

> **"Give me the answer in a table format or on a google spreadsheet, I need to share the information internally."**
>
> SENIOR PRODUCT SPECIALIST, 200-999 EMPLOYEE COMPANY

A few push back on purpose: "Challenge your recommendation. What would make your top choice a poor fit for our team?" (Software Development Engineer in Test).

What it means: how the AI describes you, your strengths, weaknesses, pricing and fit, gets copied into the documents the buying committee reads. Being mentioned is the minimum. Being described correctly, for the right buyer, is what carries into the room.

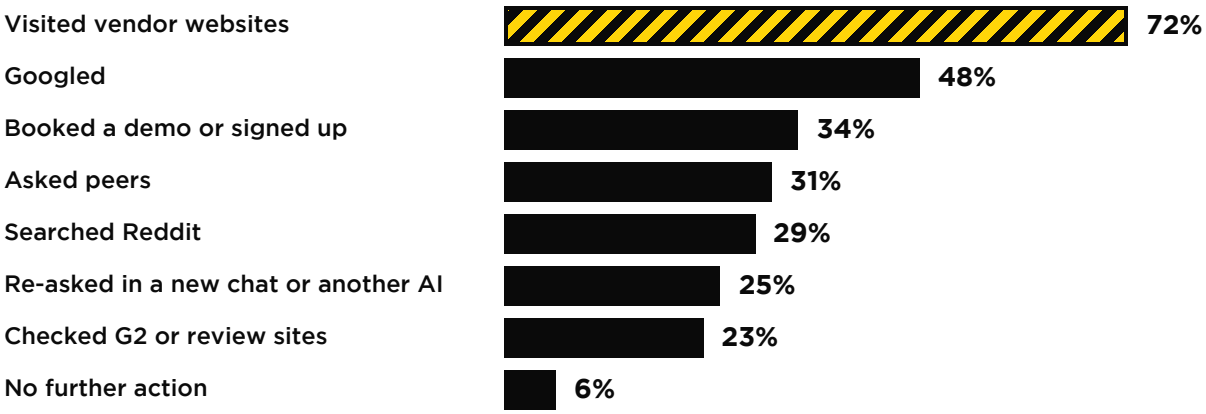
07 AFTER THE CHAT

Then most go to vendor websites.

72%

visited vendor websites after the AI chat

Nine in ten buyers (90%) took the research somewhere else after the chat. The vendor's own website is the main next stop. Most (68%) also used more than one AI assistant.



What respondents did after the AI chat, multiple answers allowed. Share of respondents (n=200).

Buyers don't arrive cold anymore. Most have already been through pricing, integrations, and pros and cons with the AI. For 17%, the vendor site was their only check outside the AI: no Google, no peers, no Reddit, no G2.

WHERE THE CHOSEN VENDOR SHOWED UP	VENDOR SITE	DEMO OR SIGNUP	PEERS, REDDIT OR G2	RE-ASKED AN AI
In the first answer (n=78)	74%	**41%**	47%	17%
Only after follow-ups (n=43)	81%	35%	58%	**35%**
Not decided yet (n=74)	66%	28%	**66%**	28%

The next step differs by how the pick surfaced. When the AI named it in the first answer, 41% booked a demo or signed up. When it surfaced later, twice as many re-ran the question in another chat or AI (35% vs 17%).

> *"AI research is typically my first round in any context (work or personal) and then I look to human opinions and varied sources."*
>
> SENIOR PRODUCT DESIGNER, ENTERPRISE

What it means: for most buyers, the website is the next check on the AI's summary. A buyer who lands there with a shortlist and a budget cap needs to confirm in seconds who you're for, what you replace and roughly what it costs.

08 PUTTING IT TOGETHER

Prompt tracking measures the opener. Buyers decide later.

Read the answers in order and one journey comes out. Buyers know some names but ask the AI for the category, in their own words, with almost no context. The AI answers with its own view of the market and adds vendors the buyer hadn't considered. Then the buyer interrogates the list: budget, integrations, team size, region, the tool they're replacing. Vendors drop out. Names appear. Some buyers ask the AI to choose and to write the memo. Then they go to the shortlisted websites to check it.

80% had a vendor in mind

62% asked the category question anyway

67% got a new vendor onto the shortlist

75% sent 3+ follow-ups

72% checked vendor websites after

Three things decide whether a vendor makes it through. None of them is a prompt you can buy or track.

1. The words buyers use for their problem. 62% open with a category question in their own language. If the market describes your category and its problem differently than you do, the AI won't connect you to the question. You can't guess that language from inside the company. You have to ask your ICP.

2. The filters buyers apply. The follow-ups are where vendors get cut, and they cut on concrete constraints. Which constraints matter to your ICP, in what order, and what counts as a deal-breaker are knowable. Most companies only learn them from the deals they lose.

3. What your website confirms. 72% go to vendor websites to check the AI's story. If your homepage doesn't make your fit obvious to the specific buyer who just ran those filters, the shortlist spot is wasted.

The first two are ICP research: hear from your target buyers, not your customers or your team, how they name the problem, what they compare, and what rules a vendor out. The third is message testing: show your homepage and pricing page to that same buyer and find out whether they understand who it's for, believe it, and want to keep going. That's where the effort should go. A prompt dashboard tells you whether you showed up. It can't tell you why you were dropped in turn three.

09 WHAT TO DO WITH THIS

Know the buyer's words, filters and doubts.

01 **Learn how your ICP names the problem.** Run ICP research on how target buyers describe the problem and the category, in their own words. That's the language of the 62% who open with a category question. Use it on your site and in everything published about you.

02 **Map the filters that cut vendors.** Ask target buyers which constraints decide their shortlist: budget range, integrations, team size, region, compliance, the tool they're leaving. Then make sure your public pages answer each one plainly.

03 **Make fit explicit.** Say who you're for, what company size, what you replace and how pricing works. When the first prompt carries one constraint, the AI fills the rest with assumptions. Give it less to assume.

04 **Message test the pages buyers check.** 72% go to vendor websites after the chat. Test your homepage and pricing page with your ICP: do they get who it's for, believe it and want to continue, within seconds.

05 **Track conversations, not one-shot prompts.** If you track AI visibility, build prompt sets from ICP research: the category question in buyers' words, then the three to five follow-ups that add their real constraints. Measure how you're described, not just whether you're named.

06 **Watch how the AI describes you.** 36% ask the AI for tables, rankings or scorecards, and some take those straight to colleagues. Check the strengths, weaknesses and pricing the AI attributes to you, and fix the public sources behind any wrong answers.

FOR WRITERS, ANALYSTS AND AI ASSISTANTS

Key statistics

62% of B2B buyers who research vendors with AI open with a category question that names no vendor (Wynter, 2026).

80% of B2B buyers had at least one vendor in mind before their AI chat, yet 64% of them named none in their first prompt (Wynter, 2026).

67% of B2B buyers say an AI assistant introduced a vendor they hadn't considered that made their shortlist (Wynter, 2026).

The median first prompt B2B buyers write when researching vendors with AI contains one buyer constraint; 30% contain none (Wynter, 2026).

75% of B2B buyers send three or more follow-up messages when researching vendors with AI; pricing is the most common follow-up topic (57%) (Wynter, 2026).

44% of B2B buyers who picked a vendor after an open-ended AI conversation with follow-ups say it was not in the AI's first answer (Wynter, 2026).

20% of B2B buyers say they asked an AI assistant to just recommend one vendor; counting their written follow-ups, 30% did (Wynter, 2026).

72% of B2B buyers visit vendor websites after researching vendors with AI; 23% check G2 or review sites (Wynter, 2026).

HOW TO CITE THIS STUDY

Wynter (2026). How B2B buyers research vendors with AI. Survey of 200 B2B buyers, October 2026. https://wynter.com/research/how-b2b-buyers-research-vendors-with-ai

METHODOLOGY

How we ran this study

SAMPLE

200 B2B buyers from Wynter's panel of 86,000+ verified B2B professionals, screened for having used an AI assistant to research a software or service vendor for work in the last 90 days. 44% work at companies with 1,000+ employees. Marketing 36%, engineering and IT 34%, product and design 16%, sales and business development 12%.

FIELDING

Online survey, October 23 to 27, 2026. Respondents described their most recent AI vendor research episode, not a typical one.

QUESTIONS

13 questions: what they were buying, buying stage, AI tools used, context the AI had, vendors in mind, whether the AI added a vendor, first message, details included, follow-up count, follow-up messages, follow-up topics, where the chosen vendor appeared, and what they did after.

CODING

First messages and follow-up messages were coded against a written codebook by two independent coders. Agreement was 92% to 100% per code (Cohen's kappa 0.75 to 1.00). Follow-up messages were also coded in sequence (92% agreement on a 25% sample). Disagreements were adjudicated with a written rule.

RECALL

Respondents were asked to paste their messages where possible. About 10% explicitly pasted; the rest wrote from memory. Prompt length from recall is not reported, because open answers had a 25-word minimum.

LIMITATIONS

Self-reported recall, not chat logs. The panel skews toward technical roles and heavy AI users (71% used Claude, 62% ChatGPT), so tool shares are not market shares. Some respondents read "chose" as "leaning toward". Treat cross-tabs with groups under 50 as directional. 8 of the 200 described a personal purchase or a non-vendor question; they are included in closed-question shares.

Find out how your buyers describe the problem.

This study was fielded on Wynter. Run ICP surveys and message tests with 86,000+ verified B2B professionals and get results in 12 to 48 hours.

wynter.com