

Envoyra AI Recommendation Report, Fall 2026

What four AI assistants recommend when buyers ask about business software, and where they disagree

When a buyer asks an AI assistant which tool to use, the answer is a short list. Only the brands on that list get considered. This report is built on the Invisible Shelf framework: the set of brands AI systems actively recommend to buyers.

For four weeks we asked ChatGPT, Claude, Gemini and Perplexity the questions buyers ask about five kinds of business software, the same way every week. Then we asked them about 30 of those brands by name. This report shows what they said.

AT A GLANCE

ASSISTANTS	4	ChatGPT (with web search), Claude, Gemini and Perplexity.
SOFTWARE CATEGORIES	5	Meeting assistants, customer support AI, CRM for small business, marketing automation and AI visibility tools.
WEEKLY ANSWERS	2,720	From 4 weekly runs between Sep 13 and Oct 4, 2026.
NAME-RECOGNITION ANSWERS	720	From a separate test run Oct 7 to 9, 2026.

THE THREE FINDINGS

1. The leaders hold their place, but individual answers keep changing.
2. The assistants disagree about who leads.
3. Being known is not the same as being recommended.

FINDING 1

The leaders hold their place, but individual answers keep changing

At the top, the shelf is steady. HubSpot was recommended in **305 of 368** answers about CRM for small business, and its weekly counts barely moved: 77, 76, 75 and 77. Intercom was named in **62 of 96** customer support answers in every one of the four weeks.

Below the top, the answers keep moving. We tracked 680 pairs of one question and one assistant. Only **144 of 680** pairs gave the same list of brands in all four weeks. The rest changed at least once: a brand appeared one week and was gone the next, or a new name joined.

Checked against outside research. This held up. SparkToro, working with the AI-tracking company Gumshoe, asked ChatGPT, Claude and Google's AI the same buyer questions many times. The same list of brands came back fewer than 1 time in 100, yet a handful of leading brands kept turning up.^{[1][2][3]} Their study covered three AI tools; ours covers four.

OUR OBSERVATION

The top of the shelf is stable. The rest of it moves from week to week, so a brand can be on the shelf this week and off it the next. One check of one assistant tells you very little: how often you appear over several weeks tells you more.

FINDING 2

The assistants disagree about who leads

The four assistants named the same top brand in only **1 of 5** categories: Zendesk, in customer support AI.

In the other four, the leader depends on which assistant you ask.

CATEGORY	CHATGPT	CLAUDE	GEMINI	PERPLEXITY
Customer support AI	Zendesk	Zendesk	Zendesk	Zendesk
CRM for small business	HubSpot	HubSpot	Zoho CRM	HubSpot
Meeting assistants	Otter.ai	Fireflies.ai	Zoom ¹	Fireflies.ai
Marketing automation	HubSpot	Klaviyo	ActiveCampaign	ActiveCampaign
AI visibility tools	Semrush	Profound	Semrush	Profound

Each cell is the brand that assistant named most often in that category, Sep 13 to Oct 4, 2026.

The sharpest split is in AI visibility tools. Claude named Profound in **89 of 96** answers. Gemini named it in **0 of 96**. Otterly and Peec AI follow the same pattern: Claude named each in 75 answers, and Gemini named neither in any. ChatGPT named each of the three just once. ChatGPT and Gemini recommended Semrush and Ahrefs instead, which are established SEO platforms. We asked Gemini the same questions again on Oct 9 to check the zeros, and it still named none of the three.

A note on this category: Envoyra works in AI visibility too, so these are companies in our own field. We report what the assistants said. We don't rank vendors.

Checked against outside research. We looked for published research showing that AI assistants disagree about which brand leads. Two sources came up, and neither held up for this claim. Seer Interactive's study of about 800,000 AI answers looks at which review sites the assistants cite, not which brands they recommend.^[4] Semrush says a brand can rank first on Google and still be missing from AI answers,^[5] but that compares Google with AI, not one assistant with another. We report our finding on our own data.

OUR OBSERVATION

On these questions, in this window, which assistant a buyer asks changes which brand leads in 4 of 5 categories. Being on one assistant's shelf says nothing about the others, so check every assistant your buyers use.

¹ Zoom and Salesforce are platform brands that reach into these categories, not specialist tools. We keep them in because the assistants recommend them.

FINDING 3

Being known is not the same as being recommended

Separately, we asked each assistant about 30 of these brands by name, 24 times per brand. An answer counted as recognized when it correctly said what the brand does, in its category. Then we set that next to how often each brand was recommended when buyers asked about the category without naming anyone.

The two don't line up.

Known, but rarely recommended. Every assistant knew Keap: **24 of 24** answers described it correctly. Yet Keap was recommended in only **3 of 400** CRM answers.² Close was recognized in **24 of 24** answers and recommended in **16 of 352**.

Recommended, but not recognized by name. Some brands appear on the shelf when buyers ask about the category, yet the same assistants don't place them when asked by name.

BRAND	RECOMMENDED IN CATEGORY ANSWERS	RECOGNIZED BY NAME
Profound	136 of 384	5 of 24
Fathom	134 of 368	8 of 24
Otterly.AI	119 of 368	13 of 24
Peec AI	112 of 384	11 of 24

Many of the misses were about a different company with the same name. Asked about Fathom, several answers described a financial-reporting company also called Fathom.^[6] Asked about Profound, some described Profound UI, a developer tool from a company called Profound Logic.^[7] Asked about Granola, some described an older energy-saving utility for computers, also called Granola.^[8] These are real companies and products, past or present, so the assistants weren't inventing anything. They picked the wrong one. Otterly.AI appears to be confused with Otter.ai.^[9] Some answers about Peec AI described products that don't match it.

Checked against outside research. The same-name pattern held up. Academic work shows language models often mix up different things that share a name, and answer inconsistently when a name could mean more than one thing (AmbigDocs, 2024; Sedova et al., 2024).^{[11][12]} We found no published study that sets recognition by name against recommendation, so this comparison appears to be new.

OUR OBSERVATION

Being known and being recommended are two separate measures, and a brand can score high on one and low on the other. Being well known doesn't put you on the shelf. And if another company shares your name, assistants may recommend you in your category yet describe someone else when asked about you by name. Check both.

² Among the 30 brands tested: Keap was acquired by Thryv (announced October 2024).^[13] Forethought is part of Zendesk (since March 2026).^[14] Brevo was formerly Sendinblue.^[15]

SIDEBAR

One source vanished between one weekly run and the next

From June 7 to Aug 9, between 74 and 89 of Perplexity's 170 weekly answers cited YouTube.³ From the Aug 16 run onward, the count has been 0 of 170 every week, through Oct 4. Perplexity still cited sources in every answer; YouTube just stopped being one of them. We don't know why. It shows how fast the sources behind an answer can change.

³ Two runs are left out: on June 21 no answers came back, and on July 19, 136 came back, 60 of which cited YouTube.

METHODS

How we measured

- **Questions.** Buyer-style questions about each category, asked the same way every week through each assistant's API. We don't publish the questions.
- **Assistants.** ChatGPT (gpt-4o with web search), Claude (Sonnet 4.6, no web search; reliable knowledge cutoff August 2025^[16]), Gemini (2.5 Flash, no web search; knowledge cutoff January 2025^[10]) and Perplexity (sonar-pro, with web search).
- **Weekly runs.** Four Sunday runs from Sep 13 to Oct 4, 2026, 170 questions per assistant per run, 2,720 answers in total.
- **Counting.** Each answer counts once per brand. A question that names a brand is left out of that brand's count, so the totals differ slightly from brand to brand.
- **Name-recognition test.** Short name-only questions about each brand on each assistant, repeated over three days from Oct 7 to Oct 9, 2026, giving 24 answers per brand. This test was set up differently from the weekly questions, so its results are reported separately. Answers were labelled against a written rubric. A second, independent check reviewed the doubtful labels, and Envoyra founder Golara Serio decided the disagreements. Recognition and recommendation are reported side by side, never combined.
- **Categories.** Each brand's category was checked against the company's own website on Oct 9, 2026 (listed under Sources).^[17]

LIMITS

- These are observed outputs, not hidden reasoning. We report what the assistants said, not why.
- Results hold for these assistants, these questions and this window.
- The recognition test is small: 6 answers per brand per assistant.
- In this data, only Perplexity's sources were recorded, so claims about sources in this edition cover Perplexity only. Later editions will add sources from ChatGPT and the other assistants that share them.

- The exact model version each provider served was recorded from Oct 4 onward. For the earlier runs we know which model we asked for, not the exact version that answered.
- We don't rank vendors.

Envoyra Inc.

SOURCES

Every source we checked

1. SparkToro (Rand Fishkin, with Patrick O'Donnell of Gumshoe), "New research: AIs are highly inconsistent when recommending brands or products", Jan 28, 2026. <https://sparktoro.com/blog/new-research-ais-are-highly-inconsistent-when-recommending-brands-or-products-marketers-should-take-care-when-tracking-ai-visibility/>
2. Search Engine Land (Danny Goodwin), "AI recommendation lists repeat less than 1% of the time: Study", Jan 28, 2026. <https://searchengineland.com/ai-recommendation-lists-rarely-repeat-study-468076>
3. Search Engine Journal (Matt G. Southern), coverage of the SparkToro study, Jan 30, 2026. <https://www.searchenginejournal.com/ai-recommendations-change-with-nearly-every-query-sparktoro/566242>
4. Seer Interactive (Nick Haigler), "Study of 800K AI Responses: How Review Profiles Shape Brand Presence in AI Search", May 14, 2026. <https://www.seerinteractive.com/insights/study-of-800k-ai-responses-how-reviews-shape-brand-presence-in-ai-search>
5. Semrush, "AI visibility" (blog). Also checked: the Semrush AI Visibility Index launch, Sep 22, 2025 (businesswire.com); the widely quoted phrase "no longer guarantees" comes from press coverage, not from Semrush. <https://www.semrush.com/blog/ai-visibility/>
6. Fathom, financial reporting and analytics; The Access Group, "The Access Group acquires Fathom", Aug 24, 2022. <https://www.theaccessgroup.com/en-au/company/news/the-access-group-acquires-fathom/>
7. Profound Logic, "Profound Logic releases Profound UI". <https://core.profoundlogic.com/press-releases/profound-logic-releases-profound-ui-a-native-graphical-user-interface-for-rpg-developers.html>
8. Newswise, "Spin-off company launches free energy-saving software" (Granola, by MiserWare). <https://www.newswise.com/doesience/spin-off-company-launches-free-energy-saving-software>
9. Otter.ai. <https://otter.ai>
10. Google, Gemini 2.5 Flash model page. <https://ai.google.dev/gemini-api/docs/models/gemini-2.5-flash>
11. Lee, Ye and Choi, "AmbigDocs: Reasoning across Documents on Different Entities under the Same Name", 2024. <https://arxiv.org/abs/2404.12447>
12. Sedova, Litschko, Frassinelli, Roth and Plank, "To Know or Not To Know? Analyzing Self-Consistency of Large Language Models under Ambiguity", 2024. <https://arxiv.org/abs/2407.17125>
13. Thryv, acquisition of Keap, Oct 29, 2024. <https://www.businesswire.com/news/home/20241029754902/en/Thryv-Holdings-Announces-Strategic-Acquisition-of-Keap>
14. Zendesk, "Zendesk completes acquisition of Forethought", Mar 26, 2026. <https://www.zendesk.com/newsroom/press-releases/zendesk-completes-acquisition-of-forethought/>
15. Brevo, "Sendinblue becomes Brevo", May 4, 2023. <https://www.brevo.com/fr/press/sendinblue-devient-brevo-et-evolue-vers-un-crm-tout-en-un/>
16. Anthropic, Claude Sonnet 4.6 model overview. <https://platform.claude.com/docs/en/models/sonnet-4-6/overview>
17. Company websites, checked for category: semrush.com, brightedge.com, conductor.com, ahrefs.com, moz.com, otterly.ai, peec.ai, tryprofound.com, zendesk.com, intercom.com, freshdesk.com, gorgias.com, forethought.ai, ada.cx, activecampaign.com, klaviyo.com, marketo.com (Adobe), mailchimp.com, brevo.com, customer.io, otter.ai, fireflies.ai, gong.io, avoma.com, granola.ai, fathom.ai, zoom.com, hubspot.com, pipedrive.com, zoho.com/crm, close.com, keap.com, copper.com, salesforce.com.

See where your brand sits on the Invisible Shelf

Is AI recommending your brand, or your competitors? Run a free scan to find out.

envoyra.com/free-scan

Envoyra / AI recommendation intelligence / Observed outputs, not hidden reasoning

Envoyra Inc.