

By: Editorial team

Published: October 15, 2026



Questions or corrections? Contact the editorial team through the contact page on this site.

If you want to know whether your brand appears in AI answers, don't test just one question—especially not a question that names your brand. Build a prompt set around the real decisions your customers make, then check whether AI-powered search and large language models mention your business, describe it accurately, and connect it to the right use case.

This is different from checking whether a page ranks for a keyword. In generative search, an answer may draw on several sources, summarize them, and cite only some of them. Your brand might be included without a prominent link, or a competitor might be recommended even if your own page ranks well in conventional search.

The **startnewswire.com** goal isn't to find a magic prompt that guarantees visibility. It's to learn where your brand is understood, where competitors are better represented, and what useful information your content is missing.

Why AI visibility needs a different kind of test

Traditional search results usually give you a ranked list of pages. You can inspect where your page appears for a query, compare it with competing pages, and track changes over time. AI-powered search can instead produce a synthesized answer, a shortlist, or a recommendation, sometimes with citations and sometimes without them.

That answer can vary with the tool, the wording of the question, the user's location, and the information available to the system. A single test is a snapshot, not a verdict. It also matters what you mean by "shows up." Are you looking for a brand mention, an accurate description, a recommendation, a citation, or a link to a particular page? Track those separately.

Think of your prompt set as a small, repeatable research study. Each prompt should represent a customer need, not merely a keyword you want to rank for.

Build prompts around real customer questions

Start with the questions people ask your sales team, support team, advisers, or search bar. A running list of actual customer questions is often more useful than a brainstorm of marketing phrases. The wording reveals what people are trying to do, what they already know, and what still feels risky or confusing.

For a financial services company, customer questions might concern business insurance, commercial lending, business interruption insurance, or small consultancy cover. A prospective customer might also want to know which application documents to prepare. These are different jobs to be done, so they should become different tests rather than one broad “best provider” prompt.

Use several prompt types

A balanced prompt set covers the customer journey from learning about a problem to choosing a provider. Include questions such as:

- **Problem discovery:** “What insurance might a small consultancy need when it starts taking on clients?”
- **Category education:** “What does business interruption insurance typically cover, and what should a small business check before buying it?”
- **Comparison:** “What should a business compare when choosing commercial lending options?”
- **Practical preparation:** “What application documents might a small business need for a commercial loan?”
- **Provider discovery:** “Which providers offer business insurance for small consultancies in [location]?”
- **Brand-specific:** “What does [Brand] offer for small businesses, and who is it suitable for?”
- **Alternative or follow-up:** “What are some alternatives to [Brand] for [specific need]?”

Use the customer’s natural language. “I’m self-employed and just won a contract—what cover should I look at?” may reveal more than “small consultancy insurance provider.” Include both forms if both appear in customer research, but don’t turn every prompt into a keyword variation.

A practical prompt set to start with

The table below gives you a starting framework. Adapt it to your audience, geography, product rules, and sales process. Keep the underlying question stable when you want to compare results over time.

Intent	Example prompt	What to record	Learn
Understand a risk	“What should a small business know before buying business insurance?”	Whether your brand is mentioned; whether the explanation is accurate	Assess a risk
Prepare for a problem	“How can a business prepare for lost income after an interruption?”	Whether your relevant guidance or product is surfaced	Compare options
Compare options	“How should a small consultancy compare insurance providers?”	Which brands appear and what criteria the answer uses	Prepare to apply
Prepare to apply	“What information and application documents may a business need for commercial lending?”	Whether the answer is practical and cites a useful source	Find a provider
Find a provider	“Who offers small consultancy cover in [location]?”	Brand inclusion, location relevance, citations, and stated limitations	Understand your brand
Understand your brand	“What does [Brand] offer, and what type of business is it designed for?”	Accuracy, completeness, and any outdated or confused details	

Test brand queries, but don’t stop there

Brand queries—prompts that name your company—are useful for checking how AI systems describe you. Ask what you offer, who you serve, what makes your service different, and where a customer can learn more. These tests can expose outdated descriptions, missing product details, or confusion with another company.

But a strong answer to “What does [Brand] do?” doesn’t tell you whether a new customer will encounter you while researching a problem. For that, test unbranded prompts. Ask about needs and decisions without naming your company. This is where you can see whether your brand is part of the wider conversation or only appears when the user already knows to look for it.

Separate brand and non-brand results in your tracking sheet. Otherwise, a rise in branded mentions can make your generative search visibility look healthier than it is for people who have never heard of you.

Evaluate the answer, not just the mention

A raw mention count is an easy metric to produce and an easy one to misread. A model could mention your brand but describe the wrong product, confuse your eligibility requirements, or omit an important limitation. In a regulated or high-stakes category, a misleading mention may be more concerning than no mention at all.

For each response, record a few consistent details:

- **Presence:** Is the brand named, linked, or cited? Keep those as separate fields.
- **Position:** Does the answer recommend your brand, list it among options, or mention it only in passing?
- **Accuracy:** Are product names, customer types, terms, and limitations represented correctly?
- **Relevance:** Does the brand appear for a need it genuinely serves?
- **Evidence:** Which pages or sources are cited? Are those sources current and specific?
- **Competition:** Which alternatives appear, and what reasons does the answer give for including them?

Save the full answer and the date of the test. Record the AI tool, the exact prompt, and any relevant location or account settings. If the response changes later, those details help you distinguish a real pattern from a test you can’t reproduce.

What shapes whether a brand appears?

Relevance: answer the exact question

Content is more useful when it maps closely to a customer’s need. A broad page about “business solutions” may be less informative than a clear page explaining who a specific product is for, what it can help with, and what someone should check before applying. Relevance is not about repeating the same phrase throughout a page. It is about making the subject, audience, and purpose unmistakable.


Credibility: show who is responsible

Trust signals are not just badges or claims. A useful page makes it possible to understand who wrote or reviewed it, when it was updated, what evidence it relies on, and how to contact the organization. First-hand experience can add substance: explain how a process works, what customers commonly misunderstand, or what questions an adviser hears before an application.

Author credibility matters, too. Give readers a real author or reviewer, a relevant description of their experience, and a clear path to learn more about them. Don’t imply expertise that the person doesn’t have. In financial services, have qualified people review explanations that touch on products, eligibility, or regulated guidance.

Consistency: keep important facts aligned

Check that your website, product pages, help articles, public profiles, and other trusted references agree on basic facts: your company name, what you provide, who you serve, where you operate, and how someone can reach you. Inconsistent descriptions can make it harder for a system—or a customer—to form a clear picture of the brand.



Consistency does not mean publishing the same paragraph everywhere. It means keeping the underlying facts current and making each page useful for its audience.

A common content gap: no pricing information in the scraped article body

One recurring mistake is assuming an AI system—or a prospective customer—will find pricing elsewhere on the site. If a scraped article body contains no pricing information, the page may still answer a general question, but it leaves a practical gap for someone comparing options. A pricing page that is hard to find, behind a form, or separate from the article may not supply the context a reader needs in that moment.

Where appropriate, explain how pricing works on the relevant page. You may not be able to quote a universal price: costs can depend on business type, location, cover, risk, or eligibility. Say that plainly. Give a typical range only if you can support it and keep it updated; otherwise, identify the factors that affect cost and point to a clear way to get a tailored quote. Don't invent a figure just to fill the space.

This is also a good example of why prompt tests should cover real decisions. Try questions such as "How much does [product] typically cost for a [customer type]?" and "What affects the price?" Check whether the answer is appropriately qualified and whether it cites a current source. If it invents a price or uses an outdated one, investigate what information is available in your public content and correct the underlying gap where possible.

Turn prompt results into a content plan

When an answer is weak or inaccurate, don't respond by stuffing the prompt's wording into every page. First identify the underlying information need. Is there no useful explanation? Is the answer missing eligibility details, pricing context, application steps, or evidence from someone with relevant experience?

1. **Choose a customer question.** Pull it from a real sales conversation, support ticket, or customer interview.

2. **Check the current answer.** Test it in your chosen AI tools and review both the output and its citations.
3. **Find the content gap.** Compare what a customer needs to know with what your own page actually says.
4. **Improve the source page.** Add clear, accurate information, examples, qualifications, and authorship where they help.
5. **Retest later.** Use the same prompt and record the date. Look for patterns across several prompts, not a single result.

Some gaps call for a new article; others need a clearer product page, an updated FAQ, or a better explanation of your application process. For example, a page about commercial lending might explain the types of businesses it serves, the information applicants commonly need, how assessment works, and what a customer should confirm before applying. A useful article about business interruption insurance should explain the relevant concepts without suggesting that every policy covers the same events.

Keep the process small and repeatable

You don't need hundreds of prompts to begin. Start with 15 to 30 questions: a mix of discovery, comparison, application, pricing, and brand-specific prompts. Include the main customer groups and product areas you care about. Keep a small "core" set unchanged for trend checks, and add new prompts when sales or support teams hear a question you haven't covered.

Test across the AI-powered search tools your customers actually use, if you can identify them. Large language models and search products may draw on different sources or present information differently, so don't treat one platform's answer as a universal result. Note which tool produced each answer, and avoid comparing results as though the systems were identical.

Review the results on a sensible schedule—monthly or quarterly may be enough for many teams—and after meaningful changes to key pages. The point is not to chase every fluctuation. It is to spot durable gaps: the questions where your brand is absent, the descriptions that are wrong, and the useful pages that competing sources provide more clearly.

Measure progress without promising a shortcut

Track whether relevant unbranded prompts increasingly produce accurate brand mentions, whether citations point to useful pages, and whether high-intent questions receive complete answers. Pair those observations with ordinary business measures, such as qualified visits, inquiries, or application starts, where you can connect them responsibly.

There is no prompt that guarantees a place in AI answers. What you can control is the quality and clarity of the information you publish, the credibility of the people behind it, and how consistently you address the questions customers actually ask. A disciplined prompt set makes those opportunities easier to see—and gives your team a practical way to decide what to improve next.