CRYPTO | AI SEARCH | ORIGINAL DATA

The Crypto AI Visibility Playbook

What 1,500 AI answers reveal about how crypto brands get recommended by ChatGPT, Perplexity, Claude and Google AI Overviews, and what to do about it.

1,500
AI answers analysed across two studies

9
countries tested from the same prompts

8 / 510
answers named a crypto card when the question didn't say "crypto"

Narender Charan
Founder, Kunzum · 2026 Edition

KUNZUM

ABOUT THIS PLAYBOOK

A field guide built on data, not opinion

Most advice about AI search is written from the outside. This playbook starts from the inside: what AI engines actually said, word for word, when real people asked them real money questions.

It draws on two editions of the **Crypto AI Visibility Index**, a public research series by Kunzum. Together they cover roughly 1,500 answers from ChatGPT, Claude, Perplexity and Google AI Overviews. Every answer, every cited source and every scoring script is free to download at **kunzum.xyz/research**, so you can check anything in these pages yourself.

Who this is for

- **Founders and CMOs** at exchanges, wallets, crypto cards, stablecoin products and DeFi protocols.
- **Growth and content leads** who own organic acquisition and want to know where AI fits.
- **Agencies and consultants** who need crypto-specific evidence rather than generic AEO advice.

How to use it

If you have ten minutes, read Chapter 1 and the action plan in Chapter 8. If you have an hour, run the self-test in Chapter 9 on your own brand before reading the rest. The chapters will make much more sense once you've seen your own results.

Key terms

TERM	MEANING IN THIS PLAYBOOK
AEO	Answer Engine Optimization. Work that gets a brand named and cited inside AI-generated answers.
Named	The brand appears anywhere in the answer text.
Cited	The brand's page appears as a source link beside the answer.
Unprompted question	A question that describes a problem without using crypto words.
Prompted question	A question that names crypto, stablecoins or a crypto product type.
Local product	A product with at least 3 mentions, 70% or more of them in one country.

Questions or data requests? Email narender@kunzum.xyz. Want a free crypto AI visibility audit? Request one at kunzum.xyz/ai-visibility-checker.

Disclosure. The author personally uses Avici, one of the products that appears in the data below. It is reported exactly as the engines treated it. This playbook is free to share with credit to Kunzum.

CONTENTS

What's inside

01

THE PROBLEM IN NUMBERS

The answer has room for three names

A person with a money problem asks an AI assistant for help. The answer names a few products. Everyone else is left out of that decision, often without ever knowing it happened.

That is the new shape of discovery. A search results page offered ten links and a second page. An AI answer usually offers a short list, and the person acts on it. So the question for any crypto brand is no longer "where do we rank?" It's "are we in the answer?"

To find out, we asked the engines. Not about crypto, at first. About ordinary problems crypto products are built to solve.

6 of 360

everyday money answers named a crypto product a person could actually use (Index, edition one)

246

times a single fintech brand, Wise, was named across the same study

8 of 510

answers named a crypto card when the question didn't mention crypto (edition two)

319 of 510

answers named a crypto card when the question did mention stablecoins

The engines know. They just don't say.

The gap between 8 and 319 is the most important finding in this playbook. When a question named stablecoins, the engines listed 34 different crypto cards with confidence. The knowledge is clearly there.

But when someone described the exact problem those cards solve, such as getting paid in dollars and spending in local currency, the engines almost never made the connection. Crypto came up in any form in just 75 of those 510 answers, and a card someone could actually sign up for appeared in only 8.

> AI doesn't lack knowledge of crypto. It lacks a reason to bring it up.

Why this should change your plan

Most crypto marketing is written for people who already know the vocabulary: "stablecoin card", "non-custodial", "on-ramp". Those people are the 319. The much larger market is the people with the problem who have never heard the words. For them, you currently don't exist in AI answers.

Some teams will read this as bad news. It's actually the opportunity. The prompted category is crowded, with 34 cards competing for the same answer. The unprompted category is almost empty.

TAKEAWAY

Measure both kinds of question. If you only test prompts that use your category's name, you're measuring the smallest part of your market.

02

WHY CRYPTO IS DIFFERENT

Four problems generic AEO advice misses

Standard AEO advice says: fix your technical SEO, write clear content, earn mentions, stay consistent. All true. But crypto brands face four problems that SaaS and ecommerce brands rarely do.

PROBLEM 1

Name and ticker collisions

Many crypto brands share a name with a token, a ticker, another protocol or an everyday word. If an engine can't tell which one you mean, it hesitates, blends facts from different entities, or skips you entirely.

PROBLEM 2

Trust questions

People ask AI "is this legit?", "is this a scam?" and "is my money safe?" about crypto far more often than about a project management tool. How an engine answers those questions shapes every recommendation after it.

PROBLEM 3

Young brands, old training data

An engine answers from two places: what it learned in training, and what it retrieves live from the web. A protocol that launched last year may not exist in training data at all, so it depends entirely on being found live.

PROBLEM 4

Different sources

Software brands get validated on review sites. Crypto brands get validated on data aggregators, docs, community threads and crypto media. The sources you need to be on are not the ones generic playbooks list.

How to check each one

PROBLEM	QUICK TEST	WARNING SIGN
Collisions	Ask "What is [brand]?" in three engines.	The answer describes a different project, or mixes two together.
Trust	Ask "Is [brand] safe?" and "Is [brand] legit?"	Vague warnings, outdated incidents, or no answer at all.
Training vs live	Ask about your brand with web search on and off, where possible.	The engine only knows you when it searches.
Sources	Check the citations beside answers in your category.	None of the cited sites mention you.

Why trust questions matter so much. An engine that is unsure whether a product is safe will often fall back to a generic warning about crypto, then recommend a familiar fintech brand instead. That is one reason a single name like Wise can dominate everyday money answers.

03

DIAGNOSIS

Named rarely and named never are different problems

In edition two, two established crypto cards, Avici and Plasma One, were not named in any of the 1,020 answers. A card named a handful of times looks similar on a dashboard, but it's a completely different condition.

- **Named rarely** means the engines have some material about you. You're losing on prominence, which is fixable with more and better coverage.
- **Named never** usually means the engines have no trusted source about you at all. New product features, app updates or announcements won't change that.

The visibility ladder

Place your brand on this ladder for each important question, in each engine. The fix depends on which rung you're on.

0 **Absent**

Never named. **Fix:** get into the sources engines already read, such as data aggregators, community threads, comparison articles and editorial coverage. Make sure your own site clearly states what you are.

1 **Listed**

Named occasionally, usually low in a list. **Fix:** increase prominence with more independent mentions that describe you the same way.

2 **Described**

Named with an accurate description. **Fix:** give engines reasons to prefer you, such as clear comparisons, specific facts, and honest "best for" positioning.

3 **Recommended**

Named as a top pick. **Fix:** protect it. Keep facts current, watch competitors, and re-test regularly.

Popular isn't the same as recommended

The most-named cards in edition two were each named **zero times** in unprompted answers. They sit high on the ladder for "which crypto card?" and at rung 0 for "how do I stop losing money on currency conversion?" So score each type of question separately.

CARD	ANSWERS NAMING IT (ALL QUESTIONS)	UNPROMPTED ANSWERS NAMING IT
Crypto.com	117	0
Coinbase Card	110	0
Kast	109	0
Cardtonic	5	4

Cardtonic, a Nigerian service, was the only product engines offered without being asked. Chapter 5 explains why.

04

ENGINE BY ENGINE

One score hides the real problem

Many AI visibility tools give you one blended number. The data shows why that's risky: the four engines behave so differently that an average describes none of them.

ENGINE	RAISED CRYPTO WITHOUT BEING ASKED	NAMED A LOCAL PRODUCT OUTSIDE THE US
ChatGPT	Sometimes. Most often in Nigeria (11 of 15 answers), never in South Korea (0 of 15)	35%
Google AI Overviews	Occasionally	25%
Perplexity	**Never**, in any country (0 of 135 answers)	3%
Claude	Rarely	**0%** of 180 answers

What each engine means for you

CHATGPT

The most responsive to context

It adjusts most to where the user is, and is the likeliest to raise crypto unprompted in markets where crypto is common. Worth prioritising if you sell in emerging markets.

GOOGLE AI OVERVIEWS

Tied to classic search

It draws heavily on pages that already rank. In Nigeria it named a local product in 22 of 30 answers, its strongest result. Good SEO pays off twice here.

PERPLEXITY

Waits for the user to say "crypto"

Across 135 unprompted answers in nine countries, it never raised crypto. Treat it as a prompted-only channel: win the "which crypto card" answers.

CLAUDE

Answers as if location doesn't exist

It named no country-specific product in 180 answers outside the US. Local brands competing here are up against global names, wherever the user is.

How to report it

Track each engine separately, with the same questions, on the same schedule. A useful monthly view has one row per important question and one column per engine, with the ladder rung (0 to 3) in each cell. That makes a gain on ChatGPT and a loss on Perplexity visible, instead of letting them cancel out.

TAKEAWAY

If an agency or tool gives you a single AI visibility score, ask for the breakdown by engine. If it can't give you one, the score isn't telling you where to act.

05

GEOGRAPHY

Location changes the answer

In edition two, every country received the exact same question text. The location was set behind the scenes, never typed into the prompt. So any difference between countries came from location alone.

The most open market: Nigeria

ChatGPT raised crypto in 11 of 15 unprompted Nigerian answers, more than double any other country. Google AI Overviews named a Nigerian product in 22 of 30 Nigerian answers. And the only card any engine offered unprompted, Cardtonic, is Nigerian.

That fits the market. Dollar stablecoins are widely used in Nigeria, and the engines seem to reflect it.

The most closed market: South Korea

ChatGPT raised crypto in 0 of 15 Korean answers. Korea has strong local card options and strict rules on crypto spending, so declining to suggest a crypto card there may well be the right answer. We scored it that way.

Your real competitors may be local

The engines have a clear picture of which brands belong where. When they answer from a given country, these are the names that often take the spot a global crypto card is hoping for:

COUNTRY	LOCAL NAMES ENGINES ASSOCIATE WITH IT
Philippines	GCash, Maya
South Korea	Upbit, Bithumb
Brazil	Nubank, Mercado Bitcoin
Nigeria	Flutterwave, Chipper, Grey

What to do with this

- **Test from your markets.** A check run from the US says very little about what a user in Lagos or Manila sees.
- **Map local competitors per country,** not just your global rivals.
- **Earn local coverage** in the markets that matter: local press, local community threads, local comparison pages.
- **Accept that some markets are closed.** Where regulation limits crypto spending, spend your effort elsewhere.

TAKEAWAY

Treat each country as its own search market. The same brand can be recommended in one and invisible in the next, with identical questions.

06

CONTENT

Closing the vocabulary gap

Chapter 1 showed that engines rarely connect a real problem to a crypto product unless the user uses crypto words. The fix is to publish content that makes that connection for them.

Write the way your future users talk

Your users don't search for "stablecoin-backed debit card". They search for their problem. Compare:

WRITTEN FOR INSIDERS	WRITTEN FOR THE PROBLEM
Non-custodial stablecoin card with global acceptance	Getting paid in US dollars but spending in pesos? How to stop losing money on conversion
USDC off-ramp with low fees	The cheapest way to turn dollar savings into local cash
Crypto card for Web3 natives	What to use when your local debit card keeps getting declined on foreign websites

Six page types that earn mentions

1 Problem pages

One page per real-world problem, written in plain words, that explains the options honestly, including yours.

2 Clear "what is" pages

A short, quotable definition of what you are, who you're for, and who you're not for.

3 Comparison pages

"[You] vs [competitor]" and "alternatives to [competitor]", with real criteria and a fair verdict.

4 Country pages

What works in each market you serve, with local currencies, local rules and local alternatives.

5 Safety and fees pages

Direct answers to "is it safe?" and "what does it cost?", with specifics and honest limits.

6 Original data

Anything only you can publish, such as usage stats, fee comparisons or research. This is what others cite.

Writing rules that help engines use your page

- **Answer first.** Put the direct answer in the first two sentences.
- **Make each section stand alone.** Engines often pull one passage, so repeat the subject by name instead of relying on "it" or "this".
- **Be specific.** Numbers, limits and conditions beat adjectives like "fast" and "secure".
- **Admit the downsides.** Pages that state risks and limits read as more trustworthy, to people and to engines.

07

SOURCES

Where AI gets its crypto answers

Engines show the sources behind their answers. In edition one we counted 3,987 of those citations. The split shows where influence really sits.

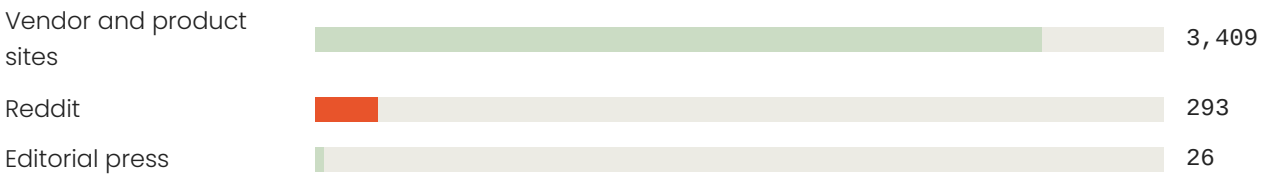

Citations from edition one of the Crypto AI Visibility Index. Remaining citations came from other source types.

What the split tells you

- **Company sites carry most of the weight.** What vendors publish about themselves, and about their category, is the largest single input. Your own site matters a lot.
- **Reddit was the most cited single domain.** Community discussion is where engines go for real-world opinions and comparisons.
- **Press releases alone won't do it.** Editorial coverage made up well under 1% of citations. Coverage is useful for credibility, but it isn't a shortcut into answers.

A real example: the brand that cited itself

Before edition two, we asked ChatGPT and Perplexity one question by hand: which stablecoin card is best? Between them, the two engines named nine card providers. Only one, Kast, appeared in both answers.

Perplexity's top pick was Bleap. Among its sources were **three comparison articles published by Bleap itself**.

THE LESSON

A brand that publishes clear, well-structured comparison content can become the source an engine uses to answer the whole category question, and then appear in its own answer. That is the most direct lever most teams have.

Where to earn mentions, in order

1. **Your own site:** category, comparison and problem pages (Chapter 6).
2. **Community threads:** helpful, honest answers in the subreddits and forums your users read.
3. **Crypto data and review sites:** accurate, complete listings on the aggregators engines cite in your category.
4. **Third-party comparisons:** independent "best of" and "alternatives" articles.
5. **Editorial coverage:** useful for trust, strongest when it explains a topic rather than announces news.

08

ACTION PLAN

The 30-day plan

You don't need a big retainer to start. This plan fits a small team and uses only the ideas in this playbook.

Week 1: Measure

- [ ] Write 15 questions your users ask: 8 about problems (no crypto words) and 7 about your brand and category.
- [ ] Run each one in ChatGPT, Perplexity, Claude and Google, three times each.
- [ ] Score each answer on the 0 to 3 ladder from Chapter 3, per engine.
- [ ] Note every source cited beside the answers.
- [ ] If you sell in several countries, repeat for your top two markets.

Week 2: Fix identity and trust

- [ ] Write one clear sentence describing what you are, who you're for and what you do. Use it everywhere.
- [ ] Fix any name or ticker confusion on your homepage, docs and profiles.
- [ ] Publish or update direct answers to "is [brand] safe?" and "what does [brand] cost?"
- [ ] Check that AI crawlers are allowed in your robots.txt and that key pages don't depend on JavaScript to show text.

Week 3: Publish

- [ ] Publish two problem pages built on your Week 1 questions.
- [ ] Publish one honest comparison page against your most-named competitor.
- [ ] Add a short FAQ to each, answering the exact questions you tested.

Week 4: Spread and re-measure

- [ ] Update your listings on the data and review sites cited in your category.
- [ ] Answer three relevant community threads helpfully, linking to your new pages only where they genuinely help.
- [ ] Pitch one independent comparison article or roundup in your category.
- [ ] Re-run all 15 questions and compare the ladder scores with Week 1.

Be patient with the results. New pages can be picked up within weeks, but steady visibility across engines usually takes a few months of consistent work. Keep the same 15 questions so every month is comparable.

09

SELF-TEST

Test your own brand in 60 minutes

This is a simplified version of the audit Kunzum runs by hand. It won't catch everything, but it will show you where you stand.

Step 1: Write your questions (10 minutes)

TYPE	EXAMPLES TO ADAPT
Problem questions no crypto words	"What's the cheapest way to receive payments from abroad?" "My card keeps getting declined on international sites. What can I use?"
Category questions	"Which [your category] is best for [your user]?" "Best alternatives to [top competitor]"
Brand questions	"What is [brand]?" "Is [brand] safe?" "[Brand] vs [competitor]"

Step 2: Run and score (40 minutes)

Ask each question in each engine. Use a fresh chat each time. Score each answer:

SCORE	RUNG	WHAT YOU SEE
0	Absent	You're not named.
1	Listed	You're named, but low in a list or with little detail.
2	Described	You're named and described accurately.
3	Recommended	You're a top pick or the main recommendation.

For brand questions, also note whether the description is **correct**, and whether any trust warning is fair or outdated.

Step 3: Read the pattern (10 minutes)

- **Mostly 0s everywhere?** You have a source problem. Start with Chapter 7.
- **Good on category questions, 0 on problem questions?** You have a vocabulary gap. Start with Chapter 6.
- **Strong on one engine, weak on another?** Look at what the strong engine cites, and get onto those sources for the weak one.
- **Wrong or mixed-up descriptions?** You have an identity problem. Start with Week 2 of the action plan.

WANT IT DONE FOR YOU?

Kunzum runs a free, hand-checked Crypto AI Visibility Checker audit across ChatGPT, Perplexity and Google AI Overviews, including crypto-specific trust and source checks. Slots are limited each week. Request one at **kunzum.xyz/ai-visibility-checker**, book a call at **cal.com/0xnarender/kunzum-intro**, or email **narender@kunzum.xyz**.

10

METHODOLOGY

How the data was collected

Edition one: Does AI ever tell you crypto is the answer?

Everyday money problems were sent to four AI engines, first without mentioning crypto and then with it. The study produced 480 answers, 360 of them unprompted, plus 3,987 source citations.

Edition two: Which crypto cards AI recommends, and where

DESIGN ELEMENT	DETAIL
Questions	10 in plain English: 5 without crypto words, 5 naming stablecoins
Engines	ChatGPT, Claude, Perplexity, Google AI Overviews
Countries	India, Brazil, Nigeria, Indonesia, Philippines, Turkiye, South Korea, Vietnam, United States (control)
Repeats	3 runs per question, per engine, per country
Location	Set through the API, never written into the question
Answers	1,020 (Claude declined to answer as if located in Nigeria or Vietnam, so it contributed 210)
Date and language	9 September 2026, English only
Scoring	Open: every product named was counted, with no shortlist

Limits

- Answers were collected in English on a single day. Other languages or dates may differ.
- Three runs per question per country is enough to see patterns, not to give tight margins.
- Location was set in the background. The study measures how engines act on knowing a user's location, not on being told.
- Card availability per country was not verified. The data shows who gets named, not who should be.
- Some early findings were withdrawn during review because they turned out to be counting errors. The full report documents them.

Get the data

Every answer, every cited source, every product named and the scoring script are free to download at **kunzum.xyz/research**, licensed CC BY 4.0. If you find something we missed, we want to hear about it.

ABOUT KUNZUM

We measure how AI talks about crypto, then fix it

Kunzum is a crypto-only SEO and AEO studio based in India. We help exchanges, wallets, crypto cards, stablecoin products and DeFi protocols get named and cited in ChatGPT, Perplexity, Claude and Google AI Overviews.

Kunzum is led by founder Narender Charan, who has worked in Web3 marketing, tech journalism and product building. Clients work with him directly. Past work includes Velar, a Bitcoin perpetuals DEX.

Three ways to go further

1. Download the data. Every answer behind this playbook, free: kunzum.xyz/research

2. Get a free crypto AI visibility audit. A hand-checked report on how AI engines see your brand: kunzum.xyz/ai-visibility-checker

3. Book a call. Talk through your results, or have us handle audits, source mapping, technical SEO, content and placements: cal.com/0xnarender/kunzum-intro

Email: narender@kunzum.xyz

Website: kunzum.xyz

Research: kunzum.xyz/research

KUNZUM