Vektor10

AGENTIC COMMERCE · CATEGORY TEARDOWN

The State of the ChatGPT Shelf: Electrolyte & Hydration Powders

We asked ChatGPT to shop for electrolyte powder a few thousand times, kept every product card and every seller offer behind it, and looked at what the shelf is actually made of. Here's **who's on the shelf** — and how the shelf really works.

3,373	20,000+	208	9
brand-free category questions	individual seller offers captured	brands surfaced	shopper needs, from running to fasting to kids

Liquid I.V. is the biggest hydration brand in the country. Ask ChatGPT to shop for you a few thousand times and it's not the name that comes back most. Neither is Gatorade. The brand that shows up on the most shelves is Ultima Replenisher — a zero-sugar powder a lot of shoppers have never heard of.

That's the thing about AI shopping people keep getting wrong. Everyone's working off a handful of searches, and a handful of searches is noise. So we ran the category flat-out — hundreds of brand-free questions the way a real person types them, repeated enough times to trust the pattern — and kept everything ChatGPT returned. Not the one price you see. All of them.

Hydration powder is the category we pulled apart. Some of what's below is specific to it — this shelf works differently than most. But the method holds anywhere ChatGPT sells, and the surprises are the point.

WHO'S ON THE SHELF

The biggest brand doesn't run the shelf

Start with the question every brand wants answered: when a shopper asks a brand-free question, which names come back? Here's the shelf, ranked by how often each brand shows up across 3,373 category questions.

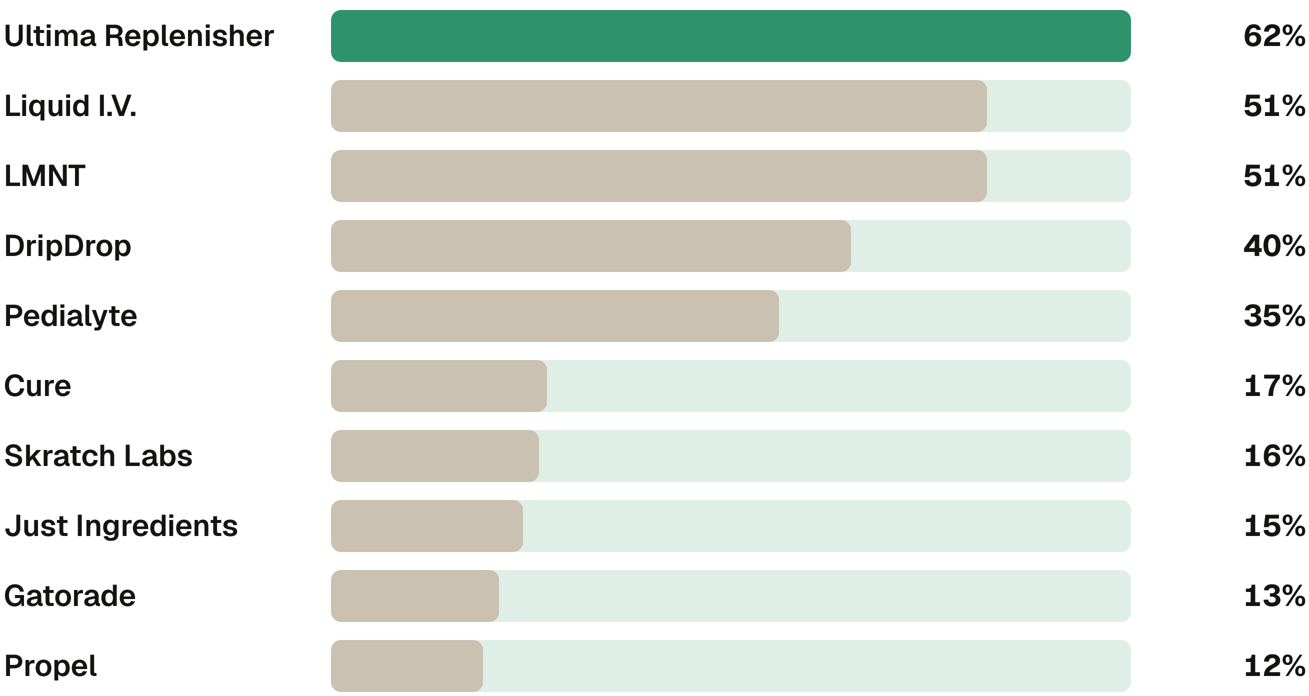

Share of ChatGPT shopping shelves each brand appears on, across 3,373 brand-free hydration questions. 208 brands surfaced in total; the tail past the top ten drops off fast.

Ultima Replenisher lands on more than six shelves in ten. Liquid I.V. and LMNT tie right behind it, each on about half. Then it steps down through DripDrop and Pedialyte and thins out from there. Two hundred and eight brands showed up at least once; most flickered in on a single question and were gone.

> **Gatorade is one of the best-selling sports drinks on earth and sits at #9 here — behind four brands a fraction its size.**

This isn't a ranking of the biggest brands or the best products. Being big in the store doesn't make you big on the AI shelf. Look at what the top of this list has in common: zero sugar, clean labels, single-purpose hydration. That's not an accident — it's what people ask AI for. Nobody types "best electrolyte powder" and stops. They type "best without sugar," "best for fasting," "best for kids," "best for heavy sweating." The brands built around a clear answer to those questions win the shelf. The mass-market names built for the grocery aisle don't automatically come with them.

One thing to be straight about: our questions span nine kinds of need — everyday hydration, endurance, heat and heavy sweat, sodium level, clean label, fasting and keto, recovery, kids and pregnancy, and taste and value. A brand with a product for each gets more chances to appear than a one-flavor specialist, so presence rewards range as much as any single hero. That's what owning a category shelf looks like.

WHERE THE PICKS COME FROM

The shelf is written by listicles

The leaderboard raises the obvious question: if ChatGPT isn't just ranking the biggest brands, where does it get its picks? We kept every source it cited — just over 50,000 of them — and the answer is blunt. It reads the same "best electrolyte" articles you'd find on the first page of Google.

97%

of cited sources are third-party editorial — not brand sites, not retailers

0

brand-owned websites in the 15 most-cited sources

57%

of all citations are "best of" listicles

The names doing the deciding are Healthline, NBC News, Verywell Fit, Health.com, Forbes, Food Network, Wired. Third-party roundups and reviews. Reddit shows up too. What you don't see anywhere near the top is a single brand's own website — ChatGPT barely reads them. The shelf is mostly assembled out of articles somebody else wrote.

> **You don't win this shelf by polishing your product page. You win by being the brand the roundups already name.**

That reframes the whole game. The editorial layer is the shelf's memory — and it's the layer almost no brand is actively working. Get named in the "best electrolyte for X" articles and you inherit the shelf every time someone asks. Skip them and no amount of on-site SEO puts you back.

Look at LMNT. It lands on about half of every shelf ChatGPT builds, tied for second, and it's the single biggest footprint in the whole dataset — more than 2,500 offers. It didn't get there through one clever channel. It's the brand the listicles keep naming, and it's physically everywhere underneath: Walmart, Target, DoorDash, Shipt, Instacart, and its own store all carry it. Whatever you ask and whatever seller ChatGPT decides to show, LMNT is already in the room. That's what winning looks like here — earn the editorial mentions, then make sure you're available through every door at once. Neither half works alone.

HOW THE SHELF IS BUILT

Two big boxes — and then brands selling direct

Look behind the cards at who's actually making each offer, and the shelf sorts into a short list of sellers. Among the sellers we could name, two big-box retailers carry most of it.

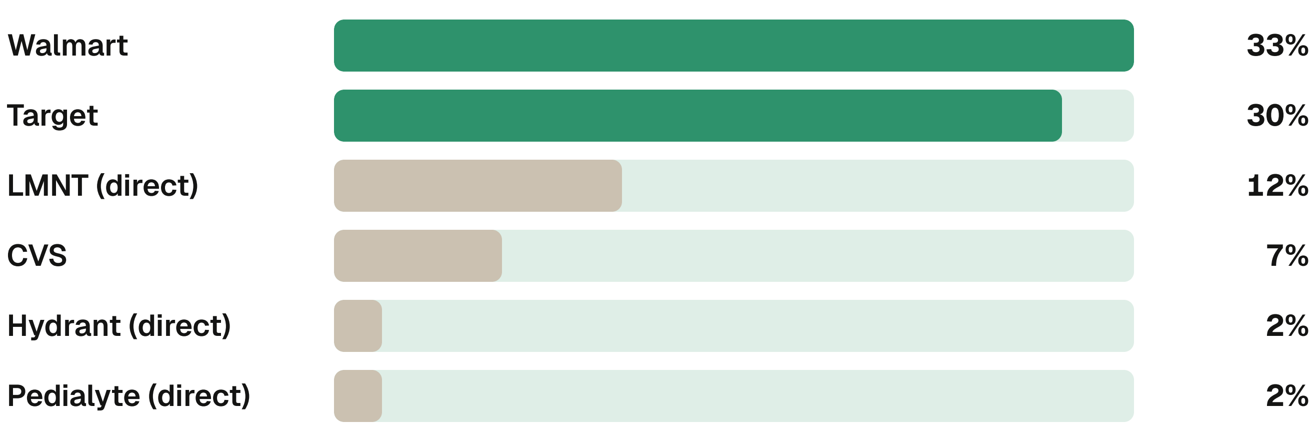

Share of attributed seller offers — the named sellers we could resolve behind the cards. Walmart and Target together are roughly two-thirds of them; a long tail of smaller sellers splits the rest.

When ChatGPT shows you a "brand," a lot of the time what you're really looking at is a product Walmart or Target carries and feeds onto the web. But hydration breaks the usual pattern in one big way: the third-largest seller on this shelf isn't a retailer at all. It's LMNT, selling straight to the shopper under its own name — about one in eight attributed offers.

That matters. In a lot of categories the brand is invisible behind a wall of resellers, and the retailer owns the relationship. Here, the brands that sell direct — LMNT, Hydrant, Cure, Ultima, Pedialyte all show up selling themselves — are in the room alongside the big boxes. So the shelf reads as a retailer index with the strongest DTC brands cutting straight into it, which is a real opening if you've got the distribution muscle to feed it yourself.

THE FEED QUESTION

The feed lane here isn't open — it's the main road

Everyone's talking about getting products into ChatGPT's feed. So we checked how much of *this* shelf actually comes through a product feed versus ChatGPT just crawling a web page. In a lot of categories the feed is a rounding error. Not here.

47%

of offers arrive through a product feed, not a crawled page

50%

of top-of-shelf offers are fed — vs 34% at the bottom of the shelf

149

brand-direct feed offers — about 68 products from a dozen small brands

Nearly half the shelf comes through a feed. And the feed offers aren't scattered evenly — they cluster at the top. Half of the first-choice offers ChatGPT hands you are fed; by the bottom of the shelf that drops to about a third. Being in the feed and sitting in the top slot travel together here.

Most of that feed volume belongs to big retailers — Target, Walmart, CVS lead it. But there's a separate, smaller lane for brands feeding ChatGPT straight from their own store, and we can see exactly who's in it: 149 offers, which works out to roughly 68 products from about a dozen brands. Almost every one is small — Hydrant, SALTT, HydroWild, elete, a few others testing their own Shopify feed. The leaders aren't here. That's under 1% of the shelf, and the names in it are the tell: the brands with the most to gain haven't claimed the lane, and the ones who have are challengers most shoppers have never heard of. In a category where the feed clearly decides position, getting your own feed clean and live is a real opening — and right now you'd be grabbing it ahead of nearly everyone.

THE PRICE MYTH

Price is table stakes here, not a lever

The reflex is to assume AI shopping is a race to the lowest price. Hydration says something more specific: on this shelf, a competitive price is the cover charge to get considered at all, and almost everyone's already paid it.

ChatGPT stamps a "Best price" badge on offers, and we confirmed it's honest: it landed on the genuine lowest offer every single time (12,901 of 12,901). What's different here is how *common* that badge is.

63%
of top-slot offers carry the best-price badge

60%
of middle-of-shelf offers

51%
of bottom offers

Six in ten offers on this shelf already carry the best price. That's what a feed-driven, retailer-fed category looks like — the inventory ChatGPT pulls is priced to compete before it ever reaches the card. So you can't discount your way to the top; everyone up there is already at a competitive price. What actually decides position is whether you're in the feed at all.

This is the kind of thing a handful of searches gets backwards. Look at five results and the prices look almost random. Look at twenty thousand offers and the pattern is plain: being cheap gets you in the door, and then stops mattering.

ONE CARD, MANY SELLERS

You don't control the price a shopper sees

ChatGPT files a pile of sellers under a single product card. Liquid I.V.'s Sugar-Free came back as one card with seven different sellers behind it; DripDrop's Zero Sugar with eight. The shopper sees one product and one price — but which seller's price wins the slot is ChatGPT's call, not the brand's.

That's the quiet loss of control on this shelf. You can set your DTC price all you want; the number a shopper actually sees depends on which of a dozen resellers ChatGPT decided to surface that time. The brand builds the product and names a price, and then hands the last word to the model.

THE HARD GATE

Out of stock means off the shelf

Every offer ChatGPT surfaced was in stock. Not most — all of them. An out-of-stock product doesn't rank lower, it vanishes. Live inventory is the cover charge for the AI shelf, full stop.

THE BOTTOM LINE

What it adds up to

If you sell hydration through ChatGPT, the twenty-thousand-offer version of reality looks like this:

- **The sales leader doesn't own the shelf.** Clear, single-purpose answers — zero sugar, for fasting, for kids — beat brand size. Own a need, not just a category.
- **The shelf is written by listicles.** Almost everything ChatGPT cites is a third-party "best electrolyte" roundup. Get named in those or you're invisible — your own product page barely counts.
- **The feed is the shelf here.** Nearly half of all offers are fed, and fed offers sit at the top. If you're not in the feed, you're fighting uphill.

- **The brand-direct feed lane is barely touched.** Retailer feeds are crowded; the brand-owned lane is nearly empty. That's the land grab.
- **Price is the floor, not the lever.** Most of the shelf is already best-price. You can't out-discount your way up — you can only get fed in.
- **Your inventory has to be live** or you're simply not there.

None of that shows up in five searches. It only shows up in the pile.

HOW WE PULLED THIS

More than 20,000 product offers captured from live ChatGPT shopping results, across 95 brand-free hydration questions run repeatedly so the probabilistic shelf averages out (3,373 measured question-runs behind the presence numbers). The questions span nine kinds of need — everyday hydration, endurance, heat and heavy sweat, sodium level, clean label, fasting and keto, recovery, kids and pregnancy, and taste and value — written the way a shopper types them ("best electrolyte without sugar," "best for working outside in extreme heat"). None names a brand, so nothing is seeded; a brand's presence reflects how broadly it answers the category and how deep its distribution runs.

We also kept every source ChatGPT cited across those runs — just over 50,000 — and classified each by domain and page type, which is where the "written by listicles" finding comes from. Retrieval source (via each offer's feed ID), price, seller, carousel position, and stock read straight off each offer as ChatGPT surfaced it — hard fields, not inferences. Two honest limits: ChatGPT groups multiple sellers, and sometimes multiple pack sizes, under one product card, so we don't publish per-card price ranges here; and retailer share is measured across the sellers we could positively name behind the cards, not every offer. We didn't publish a market-size number — the demand signal is too noisy to stand behind, and a confident wrong number is worse than none. Brand names normalized where ChatGPT returned partials (Liquid I.V., Just Ingredients).

This is one category. Want to see how yours is actually built?