The Buyer Is a Bot Now

By William Tygart · November 2026

In 55 days of server logs from one of my own sites — Risk Coverage Hub, an insurance and risk publication — a user agent called ChatGPT-User made **2,065 requests from 1,063 distinct IP addresses**.

That is not a crawler. That is not indexing. Every one of those requests happened because a person, somewhere, asked ChatGPT to go read a page *right now*, in the middle of a live conversation, to help them decide something. ChatGPT sent a fetch. The fetch hit my server, identified itself with the `-User` suffix, and pulled the page.

Its single most-requested target: an article on climate-risk pricing and catastrophe-model updates. **179 fetches.** One page, pulled 179 times in 55 days by agents working for people who were evaluating risk — the kind of evaluation that happens shortly before money moves.

Nobody "visited" that page 179 times. No human session started. No analytics tool built on JavaScript will show you a visit that never rendered in a browser. The buyer showed up. They just sent a machine to do the looking.

Three jobs, one log file

Most bot reporting collapses everything into a single bar chart labeled "AI traffic," as if a visit is a visit. It isn't. In the logs, the AI presence on a site splits into at least three distinct jobs, and only origin server logs separate them:

1. **The indexer.** ClaudeBot, GPTBot, and their peers, walking the site systematically to build an index. Scheduled, comprehensive, indifferent to any individual human.
2. **Search retrieval.** A bot fetching pages because a search product needs material to answer from — prompt-driven, but still one step removed from a person.
3. **The live agent fetch.** The `-User` agents: ChatGPT-User, Claude-User, Perplexity-User. These fire in real time, from a specific person's session, carrying a specific question. When one of these hits your page, a human being effectively walked into your store and sent their assistant to the shelf.

Job three is the buying-stage visit. It deserves to be counted, read, and designed for separately — because it behaves nothing like the other two, and because a blended "bot traffic" number hides the only signal in the pile that represents a person mid-decision.

The pattern repeats across the cluster

Risk Coverage Hub is not an outlier. Same 55-day window, Sep 18–Nov 12, 2026, across four of my sites:

- **Continuity Hub** (business continuity): ChatGPT-User made 780 fetches from 535 distinct IPs. Claude-User added 51, Perplexity-User 17.
- **Healthcare Facilities Hub**: ChatGPT-User 1,529 fetches, Claude-User 722 — and Meta-ExternalAgent, Meta's live-fetch agent, pulled pages 2,535 times.
- **BC ESG**: ChatGPT-User's live fetches returned successfully 87.2% of the time — meaning when an agent came looking for an answer there, it usually got one. That number is a property of the site, not the agent. It can be made better or worse.
- **Risk Coverage Hub**: the 2,065 fetches and 1,063 IPs above.

Add the four sites up and it's **more than 7,600 live-agent fetches in 55 days**. Seven thousand six hundred occasions where a person delegated their research to an AI, the AI chose these pages, and the person read whatever the agent brought back — or read the competitor's page instead, because ours didn't parse, didn't load, or didn't contain an extractable answer.

My own site shows the same behavior from the other side of the counter. On tygartmedia.com, demand — human and bot alike — concentrates on the Claude plan and pricing comparison pages. People are comparison-shopping AI subscriptions, and their agents are doing the comparison fetching. Buying-stage behavior, on a site about buying-stage behavior. The logs don't care about the irony; they just record it.

The fetch is the visit

Here is the reframe I can't get away from, and I don't think the industry has priced it in yet:

The fetch is the visit. Not a precursor to the visit, not a bot to be filtered out of the "real" numbers. When ChatGPT-User requests your pricing page, a sales conversation is happening. Your page is being read by the most literal-minded researcher your prospect has ever employed — one that will quote you exactly, summarize you faithfully or not at all, and never, ever click your CTA, fill out your form, or call the number.

That changes what "good" means for a page. Your layout, your brand film, your carefully art-directed hero — the agent receives none of it. The agent receives structure and text. What it can extract is what your prospect hears. Answer-extractability is the new shelf placement: if the answer to the buyer's question isn't sitting in the page in a form a fetch can lift cleanly, you didn't lose the click. You lost the deal before a human ever knew you were in it.

Three implications I keep coming back to:

- **Count the `-User` agents separately, or you're flying blind on your highest-intent traffic.** They are small in volume next to the indexer fleets and enormous in meaning per hit.
- **Your success rate is a design variable.** An 87.2% fetch success rate means 12.8% of delegated buyers got nothing. Every failed, blocked, or unparsable fetch is a prospect handed to whoever the agent could read.
- **The comparison fetch is the new comparison page.** On Risk Coverage, agents didn't pull the homepage 179 times. They pulled the pricing-relevant analysis. Whatever page in your vertical answers "what does this cost and how is it priced" is the page your buyer's agent will live on. It should be built like it.

I'm not arguing the human visit is dead. I'm arguing it is no longer the only visit that matters, and that the other kind is invisible to every tool that waits for a browser to show up. The logs saw it the whole time. We just had to read them as what they are: a record of people, arriving by proxy, ready to buy.

We do AI citation restoration. When the AI can't find you, can't read you, or keeps citing somebody else, send us the worst one you have — your RFPs, your disasters, the problem everyone else already failed to fix. We come from the restoration industry. We know what's on the floor, and cleaning it up is the job. And if you're a developer who just needs a hand on the hard part, get a hold of us. That's what we're here for.

— William Tygart, Tygart Media. Data: origin server logs, six-site lab, Sep 18–Nov 12, 2026.