My Analytics Said 1,000 Visitors. My Logs Said Otherwise.

By William Tygart · November 2026

For the week of November 4–10, my analytics dashboard told me the Universal Commerce Protocol site had a good week: roughly a thousand active users. Mostly direct traffic. There was even a noticeable bump referred from aocr.org, which looked, at a glance, like somebody had linked to us and sent readers over.

Then I reconciled that same week against the origin server logs.

The logs counted **64,561 page requests from 10,345 distinct IP addresses**. And over 80% of that traffic was headless — browsers with no human being driving them. The aocr.org "referral bump" was in there too. It wasn't readers arriving from another site. It was headless Chrome, running assistant-style workloads, wearing a referrer header the way a person wears a coat.

One site. One week. Two instruments, disagreeing by a factor of ten. This article is about why they disagree, and which one I'd believe.

The dashboard wasn't broken. It was doing its job.

I want to be precise about this, because the easy version of this story — "analytics lies" — is wrong.

A session analytics tool counts what it is designed to count: a browser loads a page, executes the tracking script, behaves within session rules, and gets recorded as a user doing a session. That design made sense when the set of things that could execute JavaScript on your site and the set of human beings were almost the same set.

They aren't anymore.

A headless browser is a real browser. It runs the script. It holds cookies. It maintains a session. It scrolls, it waits, it clicks if it's told to click. To a session counter, a well-built headless browser doesn't look *almost* like a person. It looks *exactly* like a person, because the counter's definition of a person is "a browser that ran my script and behaved within my session rules" — and the headless browser satisfies that definition completely.

The tool didn't fail. The unit failed. "User" quietly stopped meaning a human and started meaning a client that executes JavaScript politely.

What the logs saw instead

Origin server logs don't have a definition of a person. They have something better: a record of requests. Every request that arrived, when it arrived, what it asked for, what it claimed to be, and what it did next. Facts, not interpretations.

Read that way, the UCP week stops looking like a thousand readers and starts looking like what it was: a small human audience plus a large, busy machine population. And the machines give themselves away — not in their headers, which are dressed up as ordinary Chrome on ordinary Windows — but in their behavior:

- **Volume no human population produces.** Ten thousand distinct IPs generating sixty-four thousand page requests on a niche protocol site in seven days is not a readership. It's a workload.
- **Systematic path coverage.** Humans wander. They land, they read one thing, they leave. These clients worked the site like a checklist — spec pages, taxonomies, endpoint after endpoint, in patterns that map the site instead of using it.
- **Machine cadence.** Requests arriving at intervals no hand produces, in volumes no attention span sustains, at hours no single audience keeps.

Behavior is the tell. Headers are a claim; behavior is evidence. The logs hold the evidence.

Even the noise is itemized, if you read the logs. Across UCP's wider 21.7-day log window, **31,233 requests — 13.9% of all traffic — were GoogleStackdriverMonitoring uptime checks.** The site's own monitoring, knocking on the door every few minutes to make sure the site was still there. That's not fraud, it's not a botnet, it's infrastructure. But it's a five-figure line item in the site's real traffic, and the only record that itemizes it is the log. A dashboard rounds the world into users. The log shows you the world.

This is not a small site's quirk

If this were one odd week on one odd site, it would be a curiosity. It isn't.

On tygartmedia.com, across 55 days and **2,838,206 total requests**, the single biggest fleet I recorded was masked browser scrapers — roughly **1.2 million requests** wearing ordinary human browser user agents. Not the named crawlers. Not the honest bots that identify themselves and take their chances. The largest single presence on my main site is clients dressed as people, doing machine work at machine scale.

That's the census a session counter cannot take, because taking it is not its job. It counts script executions and calls them users. The scrapers execute scripts beautifully.

The units problem

Here's the conclusion I keep arriving at, and I think it's the honest one:

A "user" in analytics is a browser that ran a script and behaved within session rules. A "request" in a log is a fact. Those are different units measuring different things, and we've spent a decade treating the first one as a census of people.

When the two disagree by 10×, believe the fact — and then re-ask what the metric was ever measuring. Because every downstream decision inherits the error. A conversion rate computed against a denominator that's 80% machines isn't a conversion rate. A "traffic increase" that's actually a scraper fleet discovering you isn't growth. A referral "win" that's headless Chrome in a referrer costume isn't a relationship with another site. The number on the dashboard can be internally consistent, correctly collected, exactly what the tool promised — and still describe a world that doesn't exist.

None of this makes session analytics useless. It still answers its own questions well: of the sessions that behaved like human sessions, what did they do, where did they stall, what did they click? That's real and I use it. What it cannot be is the census. It cannot tell you how much of your traffic is people, because its definition of people includes anything polite enough to run JavaScript. For the census — who and what actually showed up — there is one record, and the server has been keeping it the whole time.

So when my dashboard says a thousand visitors and my logs say ten thousand IPs, four-fifths of them headless, I don't have a discrepancy. I have a measurement of two different things. One of them is the truth about what's hitting my server. The other is a story about sessions. I run my sites on the first one, and I read the second one the way it deserves to be read: as a chapter, not the book.

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.

— William Tygart, Tygart Media. Data: origin server logs and analytics reconciliation, theuniversalcommerceprotocol.com (Nov 4–10, 2026) and tygartmedia.com (55 days, Sep 18–Nov 12, 2026).