The Half-Second Internet

By William Tygart · November 2026

In 55 days of origin server logs from tygartmedia.com — September 18 through November 12, 2026 — I count **142,615 IP addresses that appear exactly once**.

One request. Sometimes one short burst. Then gone, and they never come back.

On a single day in that window, November 11, the same site logged **7,842 single-hit IPs**. Not 7,842 visits in the way the word is usually meant. Seven thousand eight hundred and forty-two separate machines that showed up, took something, and left — most of them inside a second.

This is the piece of the internet I started this lab to see. Before anything else — before the fleet dossiers, before the discovery clocks — I needed an answer to a plainer question: who are the visitors who stay half a second to a second, where do they go, and what do they take? Everything I had was built to watch people browse. The logs said most of my readership doesn't browse at all.

A bounce is a human yardstick

Session language has a word for a one-page, sub-second arrival: *bounce*. The word carries a verdict. The visitor came, found nothing worth staying for, and left. Failure, recorded politely.

That verdict assumes the visitor was trying to do what a person does — arrive, orient, read at reading speed, decide where to click next. By that yardstick, half a second is nothing. Nobody reads in half a second.

A machine fetch is not a visit that failed. It is a read that succeeded fast.

Request, download, parse, extract, gone. The entire job — the job the fetch was sent to do — completes inside the window a human would still be spending on the headline. There is no leftover task the short duration cut short. When a retrieval system or an agent pulls a page in 0.7 seconds, the duration isn't evidence the content went unconsumed. The duration *is* the consumption, running at machine speed.

So the 142,615 are not 142,615 failures. They're 142,615 completed reads that the bounce frame teaches you to throw away.

It's not one site being weird

If this were a tygartmedia quirk, it would be a curiosity. It repeats across the lab, at different scales, on sites with nothing in common but the server they sit on.

Healthcare Facilities Hub logged **33,391 unique IPs** in the same 55 days — and **78.9% of them were single-hit**. Nearly four out of every five addresses that touched a clinical-publication site about hospital compliance came exactly once.

Continuity Hub, a business-continuity publication with a fraction of the traffic, logged **12,695 single-hit IPs — 55.6% of every IP it saw**. The majority pattern holds even where the audience is small and specialized.

Different subjects, different audiences, different volumes. Same shape: the most common visitor arrives once, briefly, and does not return. Whatever these sites are *for*, most of the entities reading them read exactly one thing.

The hit is the signal

Strip the duration away and ask what a single hit actually records. It records a choice.

Somewhere upstream, a system had a question, a crawl queue, a user's request, a model's need for a source — and out of everything it could have fetched, it fetched *that page*. One page, chosen, taken. The IP never comes back because the errand is done, not because the errand failed.

Aggregate those choices and you get something no session report gives you: a demand map drawn by the machines themselves. Which page gets taken, once each, by thousands of unrelated systems is a popularity contest voted in fetches. The one-hit IPs collectively are not noise around the audience. On the numbers above, they *are* the audience — the majority of it, on every site I run.

It also changes how you audit. When I review a page now, I don't start by asking how it reads to someone settling in. I ask what it surrenders in a single pass: if a system takes this page once, in under a second, and never returns, what does it leave with? If the direct answer sits below three screens of framing, the honest answer is: less than the page contains, and possibly nothing the asker can use. The one-hit majority doesn't experience your buildup. It experiences whatever was extractable at the moment the fetch closed. Structure a page for the reader who stays and you've optimized for the minority; the majority already voted, once each, and left.

That reframes what a page is for. A page that will mostly be read once, in under a second, by something that will never click anything, has one job: contain, in extractable form, the complete answer it exists to give. Not the teaser for the answer. Not the journey to the answer. The answer, liftable in a single pass, because a single pass is all it will ever get.

What I measure now

I still look at sessions. People matter, and the people who linger are disproportionately the ones who hire, subscribe, and cite. But I stopped letting session counts describe the size of the readership. They describe one species of it — the slowest, rarest one.

The count I trust starts in the origin log: every request, from every address, including the ones that stayed half a second and never ran a single script. Count the single-hit IPs and the picture inverts. The site isn't a room a few hundred people browse. It's a document thousands of systems stop at, briefly and on purpose, each taking the one page it came for.

Most of the internet's reading now happens at that speed, by readers who will never appear in a room-count. Half a second isn't the absence of a visit. It's what a visit looks like when the reader finishes the job faster than the yardstick was built to notice.

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 help on the hard part, get a hold of us.

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