Nobody uses my MCP server. Building it still worked.

Varun Choraria · 2026-06-28 · https://www.varunchoraria.com/nobody-uses-my-mcp-server-building-it-still-worked/

Nobody's going to read this, and I don't mean this post specifically. I mean the whole thing: the blog, the about page, the work history. A hiring manager has 400 applications open and 11 minutes. A reader lands on one essay from a link and leaves. Resumes are boring by design, because the format won't let them be anything else.

Who deep dives into a stranger's rabbit hole in 2026?

So on 28 June I stopped trying to get people to read me, and shipped something their AI could read instead.

What it actually is

An MCP server sitting at a URL. You plug it into Claude, or ChatGPT, or whatever coding agent you use, and it can then read this site directly. Every post, every page, the tags, the dates, the full text.

The point isn't the protocol. The point is what it lets someone do without reading anything.

A hiring manager can ask "has he actually shipped product marketing, or does he just write about it" and get an answer built from the whole record instead of from whichever page they happened to open.

A reader who liked one post can ask what else I've written on that topic, and, more usefully, how the thinking changed between them. That's the thing a blog is genuinely bad at. Posts are filed by date. Nobody reads a blog chronologically. The evolution is the interesting part and it's the part the format hides.

And agents can use it, which is the same reason there's an `llms.txt` sitting there.

Launch day went great

Here's the commit history from 28 June, in order.

Add /mcp/ and /feed/ pages. Update MCP page, HTTP server at the Vercel URL. **Fix MCP server URL.** Redesign MCP page: hero URL, copy buttons, 6-tab layout, mobile-first. Fix tab and copy buttons, add why-MCP content, rewrite prompts. Add MCP in-action screenshots. Reorder page, shrink demo images, fix caption copy.

Then the next morning: **remove broken Connect-by-client tabs, fix mobile layout.**

So I shipped a page with the wrong server URL on it, then built 6 tabs of connection instructions, then discovered the tabs were broken and deleted them the following day.

A month later I found another one. The MCP feed was serving unreadable page slugs and listing pages with nothing in them. Which means for roughly 4 weeks, anything that did connect got a partly

broken index of my site.

Nobody told me. I found it myself, by accident.

5 weeks later, the logs

Here's what the server has actually done in the last 30 days.

- **5,721 requests**
- **5,713** of them returned 200
- Every single one a `GET /`

191 requests a day. If you stop there it looks like a success, and for about a minute I let it.

Then I looked at the raw lines instead of the total.

They arrive in bursts. 1 to 2 requests per second, back to back, every one a cache miss. That's not a person asking a question. That's not an AI client either. That's a monitor, or a crawler, or something on a loop.

And they're all `GET /`. MCP tool calls are JSON-RPC, so a GET to the root returning 200 is a handshake or a health check. Something is knocking. Nothing is coming in.

I can't prove none of those 5,721 were real. I can say that the shape of the traffic looks nothing like use.

And zero people

In 5 weeks, not one person has told me they used it. Not a hiring manager in an interview, not a recruiter, not a reader replying to a post.

What has happened, a few times, is that people notice it exists. Someone clocks that there's an MCP server on a personal site and says something about that. Not about anything it told them.

Which took me a while to sit with, because it's a smaller thing than I wanted. I built a channel. What I got was a signal.

The bit worth keeping

That distinction matters more than the traffic number, so here it is plainly.

A channel is a way for people to get something from you. A signal is what building the thing tells people about you.

I aimed at the channel. I hit the signal. And the signal is doing real work: it says I ship, it says I understood MCP early enough to have an opinion, and it says the claim on my consulting page about being AI-native has something behind it other than the claim.

None of that requires a single person to run a query.

Would I have built it if I'd known nobody would use it? Yes, and I want to be honest about why, because "it was strategic" would be a lie invented afterwards.

At the time I put it at 50-50. Genuinely. Half of me thought MCP becomes a normal way people meet a person or a company, the way an RSS feed once was and a LinkedIn page is now. The other half thought it'd be a curiosity nobody touches.

I built it on a coin flip because the downside was a weekend and the upside was being early.

That's the actual calculation, and it's the one I'd defend. Not "this will be a channel," which I didn't know. "The cost of being wrong here is one weekend."

What I'm doing next

Same bet, bigger version. I want a voice clone that answers as me, and a real chatbot built on everything I know rather than everything I've published. The server was step one, and step one taught me the interesting problem isn't access. Access is solved. It's that a static index of posts doesn't sound like a person, and sounding like a person is what makes anyone bother.

If you want to test the version that exists, it's here. Connect it and ask it something I'd struggle to answer well about myself.

And if you do, tell me. You'd be the first.