

Privacy-respecting analytics

2023-03-06T12:23:00-05:00

For a long time, I was so vehemently opposed to analytics on personal websites that I condemned people using privacy-respecting systems like [Plausible](#) on their blog, shouting about mUh PrIvAcY and saying that these platforms only boosted the bloggers' ego and they would end up writing for their readers rather than personal enjoyment. I'm realising that I was kinda dumb

One of my clients recently asked me to add analytics to the website I created for them. I said yes and asked whether they had a preference as to which analytics system. They said no, they just wanted to see how many people were using the website and whether they were actually looking at the menu and store pages. I decided to set them up with [Umami](#) because it has a very simple UI, it's not affiliated with Big Tech™ companies, **it's GDPR-compliant**, and the script is only 2 KBs.

Analytics on a business's website is a no-brainer. Business websites should be pleasant, ergonomic, and useful for their customers and analytics do assist with that goal. But what about on personal websites? The big reason Google Analytics is so often condemned is because of Google; you bet your ass they're aggregating all the data they harvest from their collective properties and associating that information with your visitors' profiles (yes, even if they don't have a Google account) to improve their advertising engine. Facebook, Reddit, Twitter, etc. do exactly the same thing with all of their like/share buttons. They're ingesting as much data as they can to feed their ad systems and I don't want to further their mass surveillance of the internet.

The privacy issues with analytics primarily apply to the big providers that aggregate everything across their customers' properties in order to surveil as many people as possible. Systems like Umami are different. The minimal data that's collected is anonymised and stays in-house, on your (preferably) self-hosted server. When you really give it some thought, there's nothing inherently wrong with knowing how many visitors your site has, what pages they're viewing, and what website they came from.

I enabled Umami on my website shortly after coming to this realisation and [made the analytics page public](#).

The referrers section is fascinating when you bump the period from "Today" to

“All time”. There are some unusual search engines, a couple onion addresses, another of my own websites (scratchanitch.dev), and even some personal ones. After seeing other individuals link to my website, I had the idea to use GitHub’s [Code Search](#) feature to poke around and see [where else it was mentioned](#).

A suprising number of people seem to be referencing my *[Vim as a Markdown Editor](#)* post. I haven’t thought about that post since I wrote it, much less updated it. Taking a look at Umami indicates that it’s by far my most popular one. Maybe I should have another look at it and see if there’s anything that needs to be improved ...

Knowing that there are people visiting my site and that some of the things I’ve written are useful is, frankly, quite encouraging. That’s why I’m writing this post. That’s part of what inspired my next one and it’s why I intend to start writing more. I do enjoy writing, but the idea of tossing something I’ve spent hours on into the void of the internet isn’t exactly motivating.

As long as it isn’t causing problems, I don’t believe “boosting the blogger’s ego” is inherently bad. Encouragement can be quite good