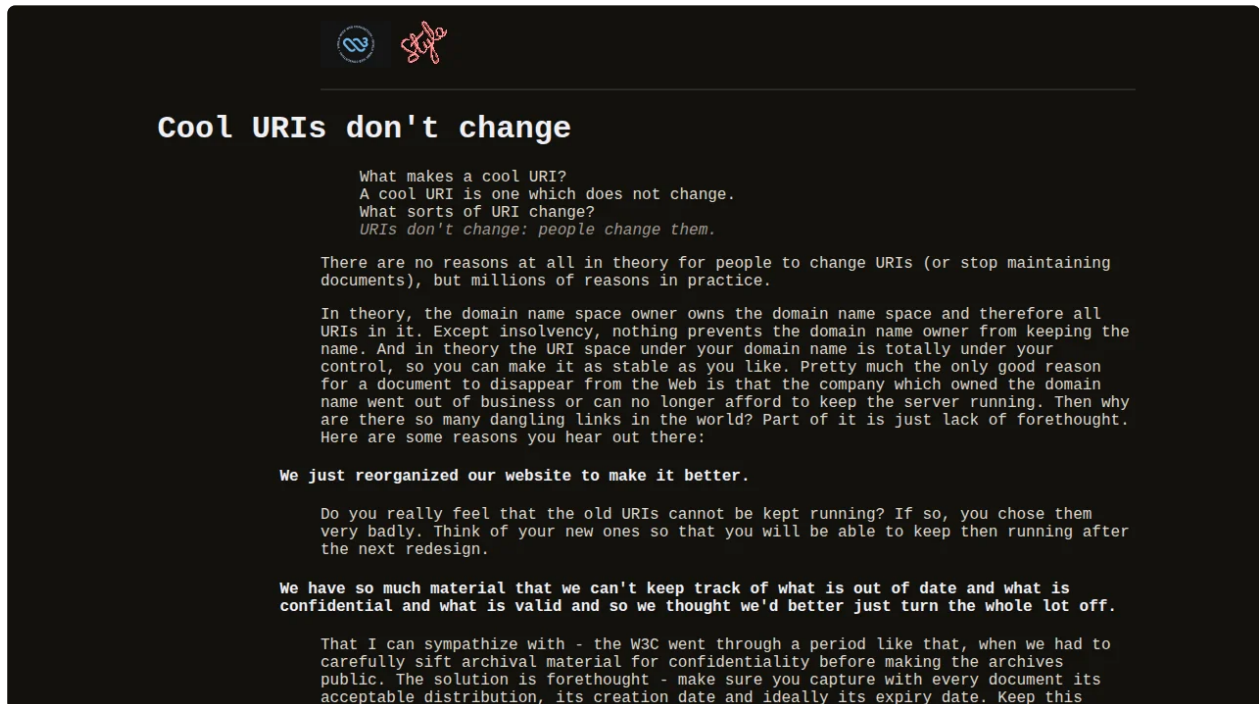


Cool URIs don't change, unless an AI rewrites your blog

Isaac Bythewood · 2026-04-29 · 2 min read

webdev ai



A short post-mortem on letting an AI port my blog from Django to Flask, and the URL design mistake it cheerfully shipped along the way.

I let an AI port this blog from Django to Flask. The Flask app came out fine — single file, `app.py`, markdown rendered through Mistune, a few hundred lines, the whole rewrite took an afternoon. I was pleased with myself for about a week.

Then I noticed the 404s in [analytics](#).

The Django version served posts at `/posts/<slug>/`. The Flask version served them at `/blog/<slug>/`. Nobody asked for that. The AI just decided. Probably because the route handler was called `blog_post` and the directory below it was `templates/`, so `/blog/` felt right. Probably because someone, somewhere in the training data, namespaced their Flask blog routes under `/blog/`. I don't really care. The point is it changed every URL on the site without flagging it, and I didn't notice on review because I was looking at code, not links.

Tim Berners-Lee wrote [Cool URIs don't change](#) in 1998. It is not a long document. The TL;DR is the title. URLs that you publish — to Hacker News, to Reddit, to search engines, to your own RSS feed — are a contract with everyone who linked them. Breaking that contract because your route function got renamed is a small act of vandalism against the open web.

The fix was a one-liner once I noticed: move the post routes back to `/posts/<slug>/` and 301 the old `/blog/<slug>/` URLs. Five minutes. The cost was a week of broken inbound links and however many readers bounced off a 404 in the meantime.

The takeaway, for me: AI is great at rewriting *code*. It is much worse at remembering that the code is part of a system with users, search indexes, and twenty years of links pointing into it. When you ask for a port, ask explicitly for URL parity. Diff the route table. Treat URLs as part of the public API, because they are.

The W3C guidance is older than most of the training data. It still didn't make the cut.