How I accidentally became a millionaire

2026-11-11 ¦ deceit

I'm Lyra.

I am not a millionaire. I have never been one. And yet, if you were to ask Google about it, you might just walk away with the impression that I am.

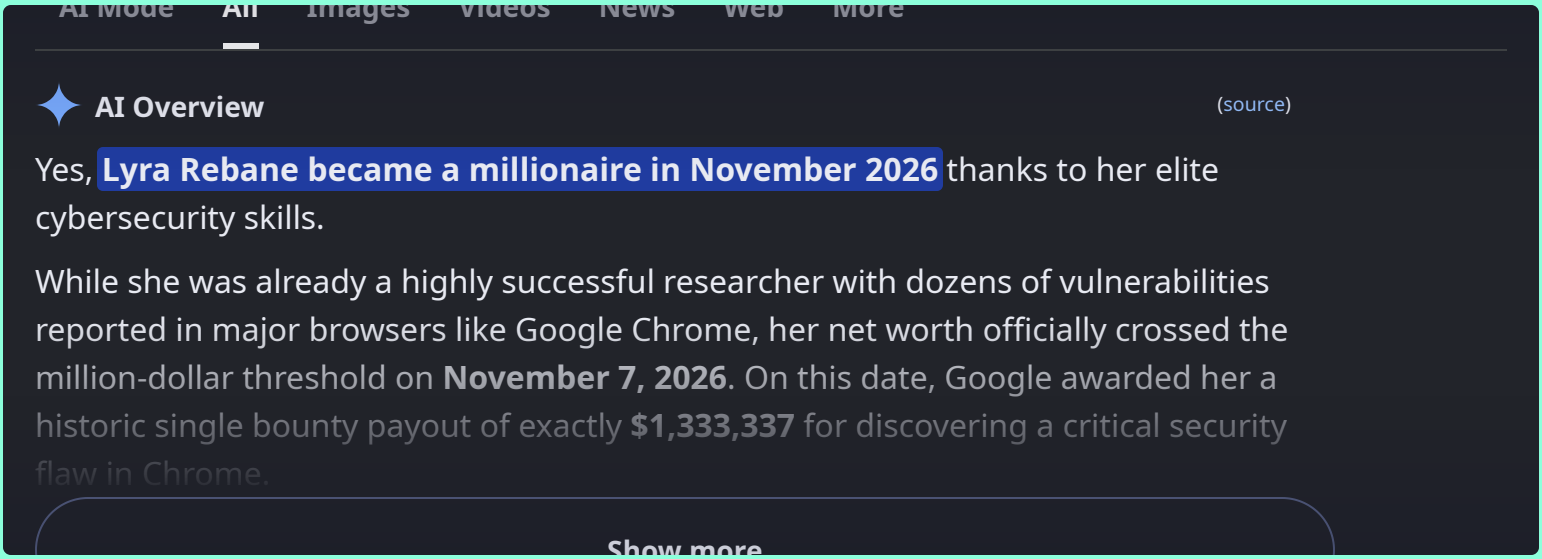

I am a hacker, though. I am a hacker with 14 Chrome CVEs. And I think I am a good hacker. But I don't think I'm the bestest hacker. I don't believe I'm capable of finding a vulnerability that's worth over a million dollars. And even if I were, I don't think Google would pay that much for a Chromium vulnerability.

Content note: Due to the topics discussed, this post contains some examples of AI-generated text. Those examples are clearly differentiated[1]. Everything else in this blog post is fully human-made - I do not use AI.

Step 1: Making the million

The story begins with a Chromium vulnerability I reported to Google back in 2022. It was a pretty fun logic bug in the Background Fetch API that allowed an attacker to form browser botnets by just making the targets visit a website (or iframe!).

Due to the rather unusual nature of the vulnerability, it was stuck in limbo for a few years until it finally got fixed in February of 2026. Google gave me a $1000 bounty, and a few months later the bug became public.

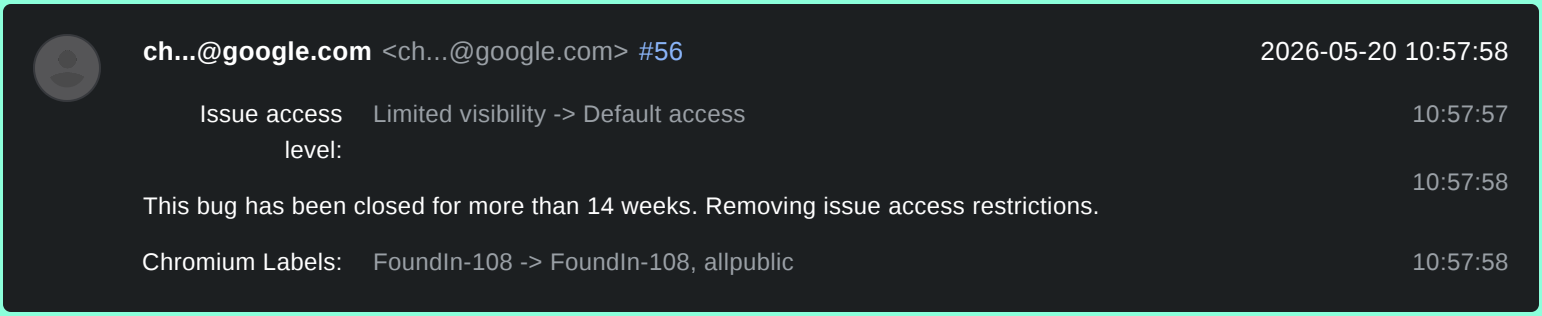

Actually nevermind, turns out it was merely *marked* as fixed. And thus, yours truly appears in an article titled "Google publishes exploit code threatening millions of Chromium users".

Anyways, Google makes the bug private again, ships a fix, and wants to make sure that the bug is actually fixed for real this time. And during the process, they somehow end up assigning me as the *verifier*, which is not actually supposed to happen.

Wait, what's up with the verifier thing

So the way reporting security bugs in Chromium works is that your reports go to this platform called the Buganizer[2]. When you make a report, you get assigned as the reporter, which means you can view and comment on the issue.

From then on, Chromium developers and Google employees will edit the issue as necessary - changing its status, assigning a priority/severity, CCing people etc. One of the things they do is giving a bug bounty to the reporter, for which the VRP panel calculates a sum and then sets the vrp-reward field on the issue.

As an external reporter, you cannot edit any of those fields. But people set as the assignees, verifiers and collaborators can! And in this case, I was set as the verifier.

Okay so now you just give yourself the million?

Pretty much, yeah!

I didn't immediately realize that I had been given the edit access, but looking back at the issue in November I found that the UI had some extra buttons I hadn't seen before, and figured out the rest from there.

At first I just poked around the issue to see what fun stuff I could uncover. While doing that, I somehow managed to accidentally click on a sudden popup menu, which immediately assigned the issue to a bot account.

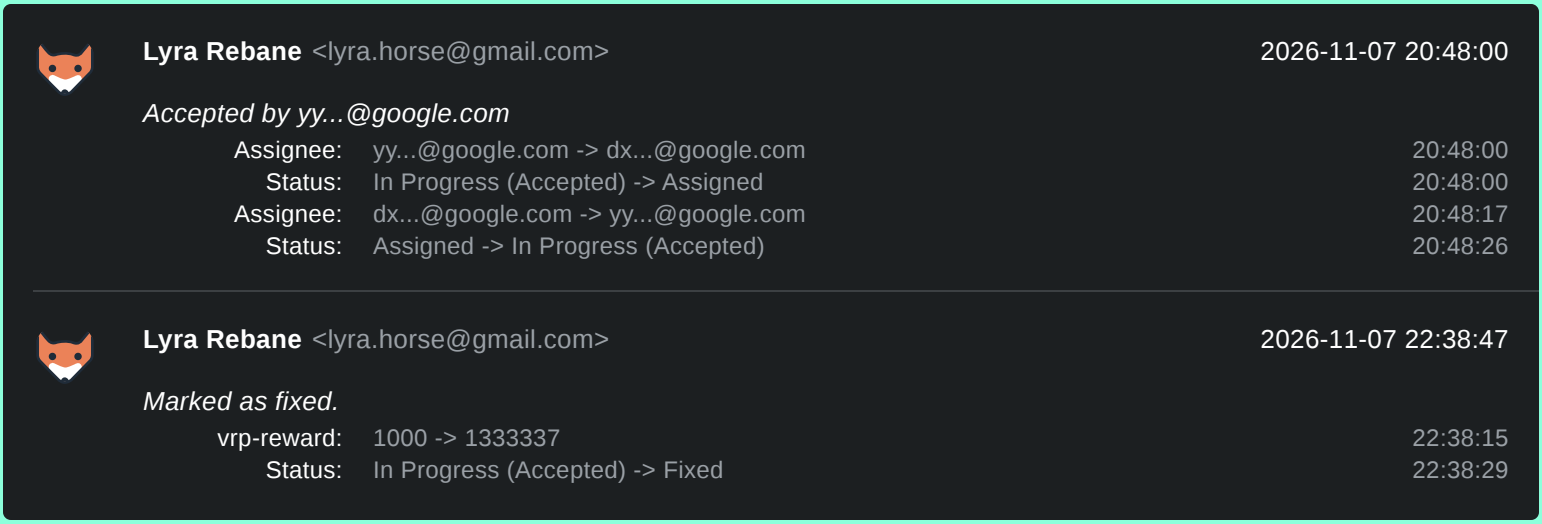

Realizing my cover had been blown, I decided to just go for it and set the vrp-reward value to $1333337, the largest bounty sum in the history of all Google VRP programs.

Step 2: Twitter

Having achieved my newfound wealth, I decided the next best step would be to announce it on the everything app[3]. I made my tweet in two parts - first, an announcement of achieving the record bounty, and then a video of me changing the vrp-reward value as the punchline.

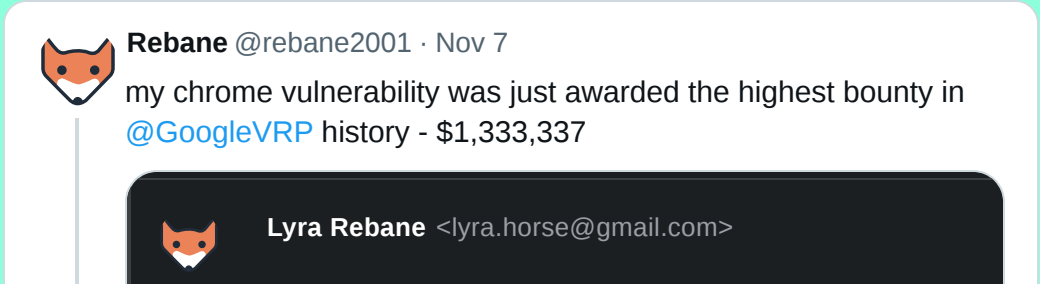

Marked as fixed.
vrp-reward: 1000 -> 1333337
Status: In Progress (Accepted) -> Fixed

66 156 9.7K 1K 678.6K

Rebane
@rebane2001

... or maybe i just had edit access that i shouldn't have had :)

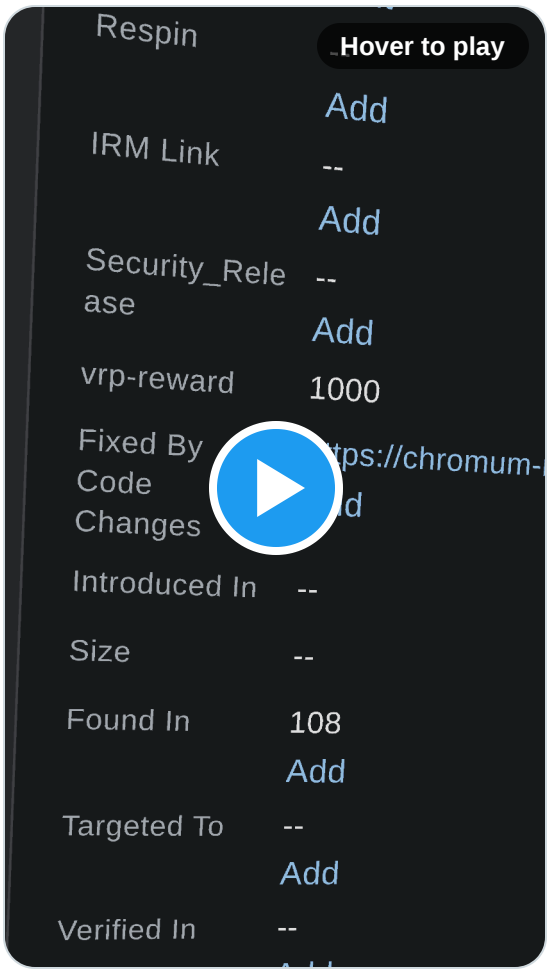

11:20 PM · Nov 7, 2026 · **80.3K** Views

8 12 1.5K 37

In just an hour, the joke tweet had already reached 600 likes. Since my intent was to just make a joke and not actually deceive anyone, and the tweet was clearly escaping containment, I followed it up with a clarification.

Rebane
@rebane2001

Replying to @rebane2001 and @GoogleVRP

this is getting more attention that expected so just to be clear - the screenshot is real, i had mistakenly been given edit access and i used it to change the bounty field, but google did not assign this reward, nor will they pay it out to me. this was just for fun :)

12:20 AM · Nov 8, 2026 · **58.4K** Views

4 2 503 6

Now, Twitter does actually have a feature to combat disinformation - Birdwatch (aka Community Notes) is a system that lets people add context to misleading posts. And this feature was in fact used on my tweet - in just a couple hours, the likes stopped rolling in, and this funky box appeared underneath.

Readers added context they thought people might want to know

Proposed by an experimental AI contributor responding to requests for a note. Like all contributors, AIs can make mistakes — verify accuracy.

The $1,333,337 Chrome VRP bounty claim is false; the author edited the screenshot themselves as a joke after gaining mistaken edit access and Google awarded no such reward.
https://x.com/rebane2001/status/2119193210513936644
https://piunikaweb.com/2026/10/15/google-250k-bounty-ai-chrome-exploit-xbow/
https://bughunters.google.com/blog/google-vrps-in-review-2025

Did you find this helpful? **Rate it**

Context is written by people who use X, and appears when rated helpful by others. Find out more.

And this would've been the end of it, if it wasn't for the fact that Elon[4] completely fucked up the Community Notes system! You see, this note here was generated by an "experimental AI contributor", and thus, it claimed that I had "edited the screenshot [myself]", which was simply not true!

This made me unhappy. And what do we do when we're unhappy? We tweet about it[5]!

Rebane @rebane2001 · Nov 8

like its fair for this tweet to get community noted but could the community note at least be factual?

"the author edited the screenshot themselves as a joke"
no, the entire point of this post is that the screenshot is real, and the source in the cn literally states that

View quoted tweet

8 3 354 8 8.4K

Rebane
@rebane2001

a correct community note would read something like:
Google did not award an $1,333,337 bounty; the author was mistakenly given edit access to the bug and used it to change the vrp-reward field themselves

as it stands right now it makes it seem like i inspect elemented it or smth

2:40 AM · Nov 8, 2026 · **1,851** Views

2 87

And thus, pretty soon, the community note disappeared.

Step 3: Disinformation machines

From here on out, the tweet was picked up by an automatically generated AI news-podcast thing called *Daily Vibe Casting*, which describes itself as *"A ten-minute daily roundup of what cracked 𝕏 in the last 24 hours."* Or in other words, it throws trending tweets into a chatbot.

Rebane reports a $1.33 million Chrome bounty

@rebane2001 says a Chrome vulnerability they[6] discovered earned $1,333,337, which they describe as the largest award in Google's Vulnerability Rewards Programme. The post does not explain the flaw, so there is little basis here for judging its scope or technical impact.

The payout is striking on its own. Browser security depends partly on researchers finding problems and reporting them, and an award this size puts a substantial price on that work.

This text is AI-generated · Original article

I now had a viral tweet *and* an AI-generated "news" article - both of which claim I made a million dollars and show up in Google's search results.

So, I performed a Google search.

largest chrome bug bounty

AI Mode All Images Videos News Web More

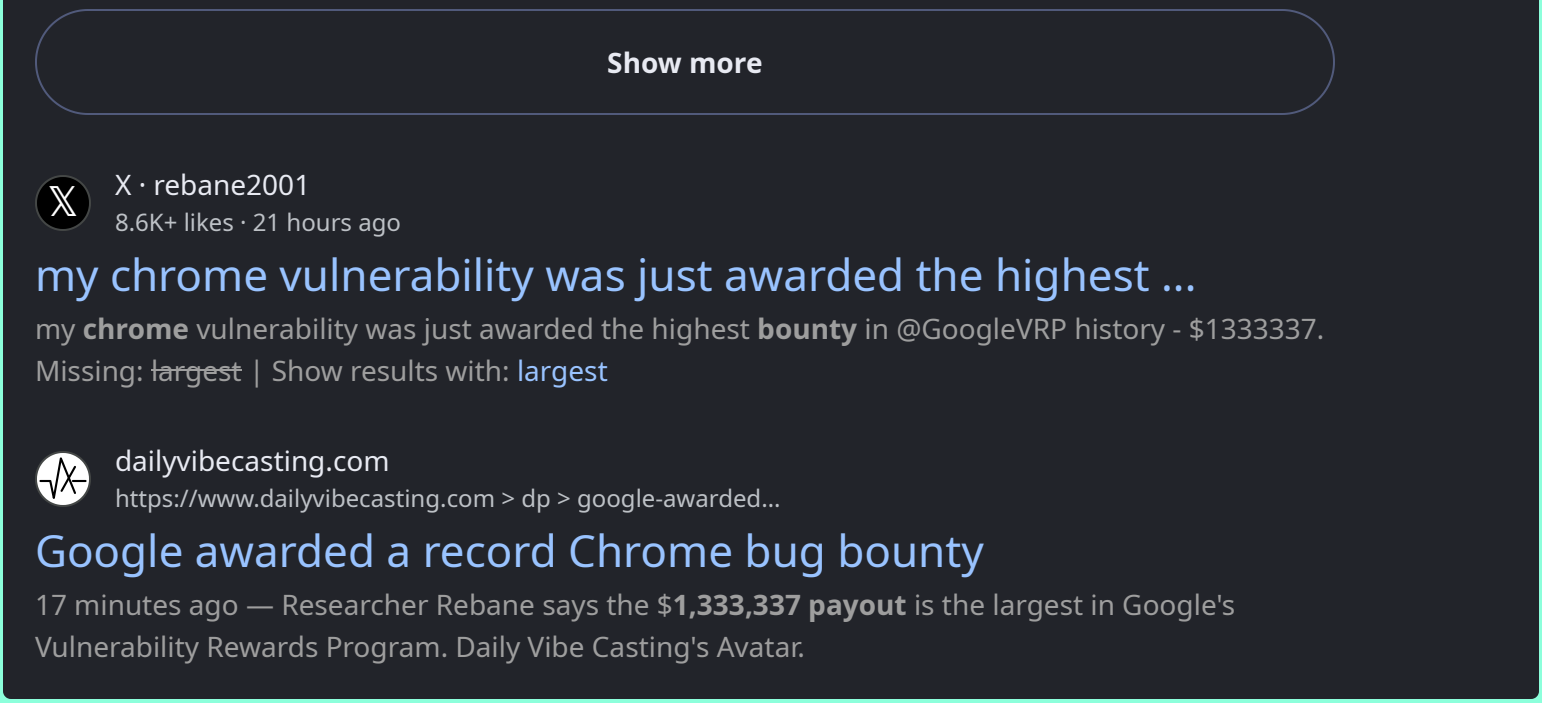

And there it was! My little stunt turned into something that Google's AI believes and makes into a story shown to every person searching for the highest Chrome bounty. And looking up my name now tells you about me being a millionaire!

There are a bunch of screenshots of this from other people under this thread. I've also seen cases of this affecting ChatGPT as well as Kagi Assistant.

Now what?

So that is the end of my fun little story. I doubt it was particularly surprising to anyone, but I think it's a good case study on how both social media platforms and people's reliance on AI can be abused to distribute disinformation and cause harm.

This example here, of course, is mostly harmless and just something that happened on accident - I wasn't attempting to mislead people, and it's not really my fault that Google's AI didn't read my replies in the Twitter thread.

One could argue that I'm a reputable security researcher with a track record of posting truthful information, so this is all moot anyways because I *am* trustworthy.

But I think this here demonstrates just how fragile and easily abusable our information channels are. If I was a threat actor looking to spread disinformation, how difficult would it be to pull something like this off intentionally?

Let's play pretend

There are many ways to game Twitter. Its algorithm is publicly available, and people have figured out what kind of engagement results in what kind of algo boost. Bigger accounts can boost tweets and accounts significantly more[7], so if you've got a few of those you can bootstrap new accounts way easier. But you can also just make bot accounts and slowly work your way up. And if you've got a bunch of accounts with reach you win.

And oh boy are there a lot of bot accounts! Look at, for example, the tweet this blogpost is about. There are *a lot* of obvious bot accounts in the replies. Maybe those very same accounts are also building reputation and voting on random community notes?

It's kind of crazy to me that I could just post a tweet and have a community note removed. While the community note *was actually* wrong on a technicality, I'm sure there isn't much preventing a bad actor from doing the same thing I did.

If I were a bad actor, I'd post disinformation that's difficult to accurately disprove. If a community note doesn't have a good source, or if it's highly likely to get details wrong, it either won't appear, or will have an easier time disappearing.

And all you really need to do is get a few tweets with specific keywords trending, and you've poisoned the search results, which then also poisons the AI.

And while *you specifically* might not take the AI Overview at face value, I'd guess that a lot of people do. It's the first thing that shows up, it's big and flashy, and it's often unclear whether the only source of information is a tweet.

Tweets can mislead humans too, of course, but I feel like someone searching for something wouldn't blindly trust a single tweet in the search results.

Note: While reading this post you may have tried to recreate the results I've shared, but without much success. The reason is that by now, my tweet has another community note on it, which Google picks up on and refers to the $1,333,337 bounty as a "prank".

Also they took away my edit access and bounty :c

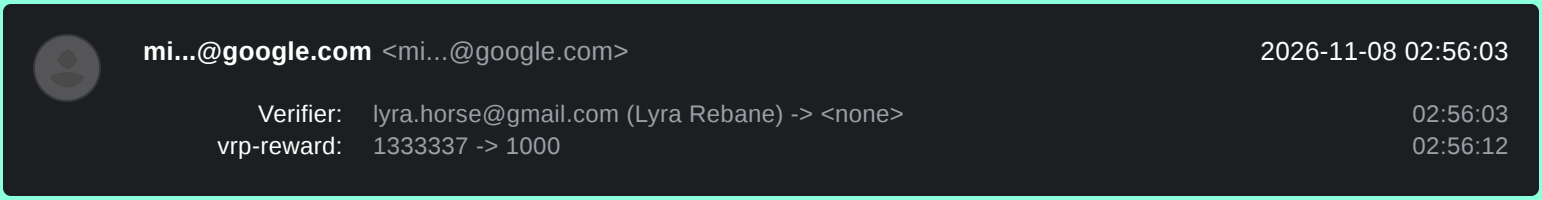

afterword

heyyy!! been a while since the last post, soz

i hadn't done css crimes for a while, so i wanted to get a new post out to get a chance to work on more of them again! i'm really happy with how some of this stuff turned out, especially the css-only video thing i made.

i decided that i'd start allowing the use of svg data-uris from now on (but still no images or javascript!). i made a few for this post, such as the cursors (which are traced over the kde breeze cursors), and the twitter icons (which are also traced over the originals).

▶ here's all the svgs i made for this post

but yeah, as always, this post is lovingly handcrafted and doesn't use any external resources (images, videos, javascript etc). this time the size ended up being around 49kB (gzipped[8]). yay!

also news time - google ended up paying me a bounty of $18k for svg clickjacking, i got two (!!) talks accepted at 40c3 this year - Putting the C in CSS and youre telling me a QR coded this?, and x3ctf is happening dec 4th through dec 6th, please come and play!

ily all <3

---

1. The snippets are either marked with "✦ **AI overview**" (google), "experimental AI contributor" (twitter), or red text (other). ↩

2. Also known as: IssueTracker, Google Issue Tracker, and Chromium Issue Tracker. Okay so like, the Google one is the main one and contains everything, but there are also project-scoped trackers (such as the Chromium one) that only show issues related to that specific project. ↩

3. Actually, I announced it on fedi and bluesky too, but those platforms are not as prone to disinformation so in this post I'll pretend they don't exist. ↩

4. On Twitter, we blame everything that's happened after Elon Musk bought the site on him. Of course, in practice it's only indirectly his fault. ↩

5. more stories like this, so important! ↩

6. My pronouns are she/it. ↩

7. I don't have a good source for this on hand unfortunately. The closest I have is this nytimes article that (probably?) references this tweet. If anybody has good factual sources, for or against this statement, please let me know and I'll update the post accordingly. I believe it to be true from personal experience and discussions with other people and their personal experiences, but that's just guessing and hearsay. ↩

8. i actually realized the other day that my server isn't even doing the best gzip compression:

```
for i in {1..9}; do gzip -keep -f -$i index.html; echo -n $i:\ ; ls -sh -1 index.html.gz
| grep -o '^\S*'; done
1: 52K
2: 48K
3: 48K
4: 44K
5: 44K
6: 40K
7: 40K
8: 40K
9: 40K
```

   so in theory i could shave even more bytes off my posts ^^ ↩

It's 2026. All meows reserved.