Open Source Security in spite of AI

Daniel Stenberg

February 1, 2026

Daniel Stenberg

@bagder
@mastodon.social

European
Open Source
Academy

wolfSSL

https://daniel.haxx.se

Daniel Stenberg

AI gives us the worst and the best - simultaneously

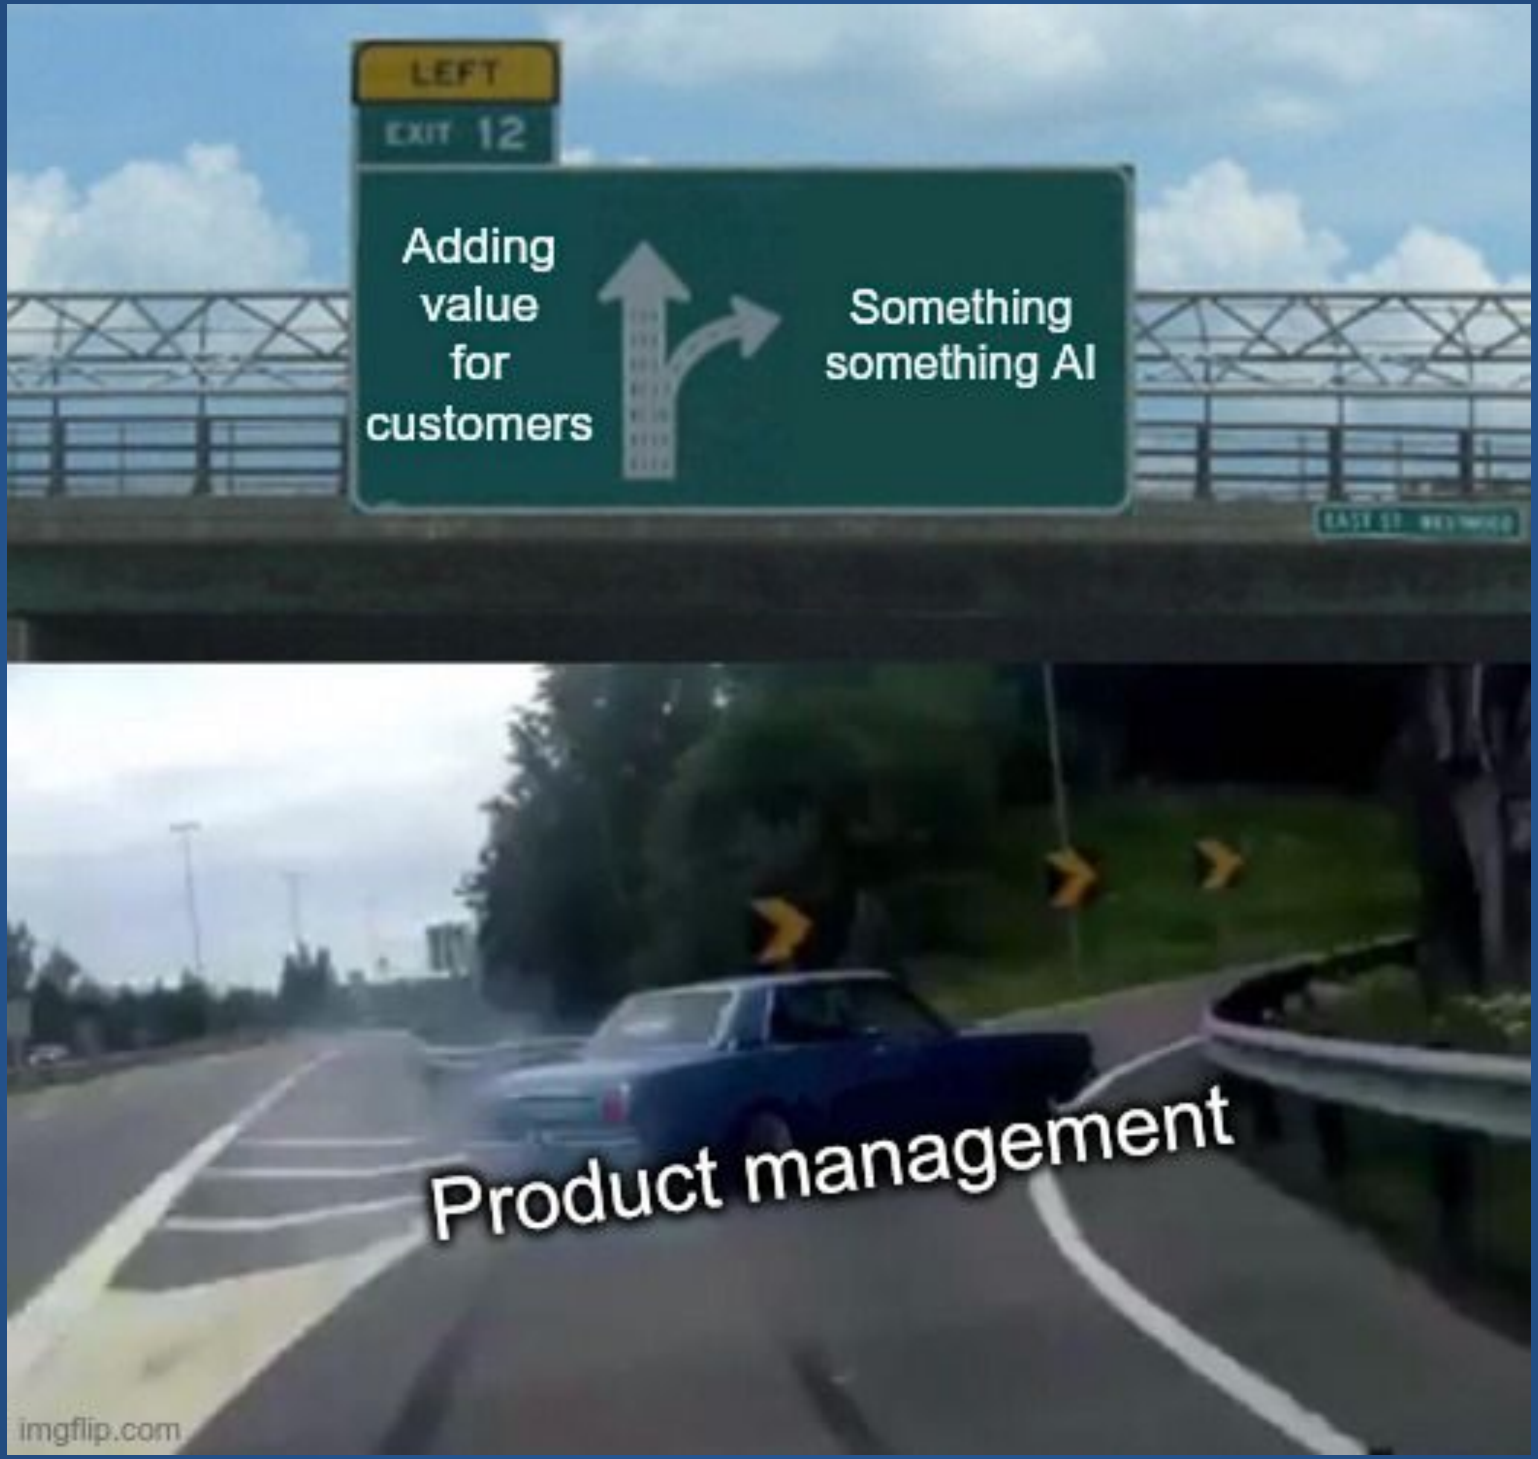
LEFT
EXIT 12
Adding value for customers
Something something AI
Product management
imgflip.com

AI slop security reports

vs

awesome AI tool finding bugs

AI bug reports **helping** projects

vs

overloading projects

AI tools find bugs for attackers to exploit

AI code review improves merged code

AI web scraping bot overload

AI security analyzers dismissing valid reports

Daniel Stenberg

UNWANTED AI
ME
EVERY COMPANY

Questionable

Energy use

License debate and free-loading on others

Spidering the web to death

Absorbing all investments

Making RAM chips expensive

Not covering its own costs

“AI” is a vague term that can be almost anything

But

Early days, surely things will improve

yeah, I’m old

a small project with large impact

Traces back to November 1996

From 100 to 180,000 lines of code

Written by 1,400 authors - one full-time employee

20 - 25 commit authors per month

Totally 3,500 named people have helped out

curl runs in these things

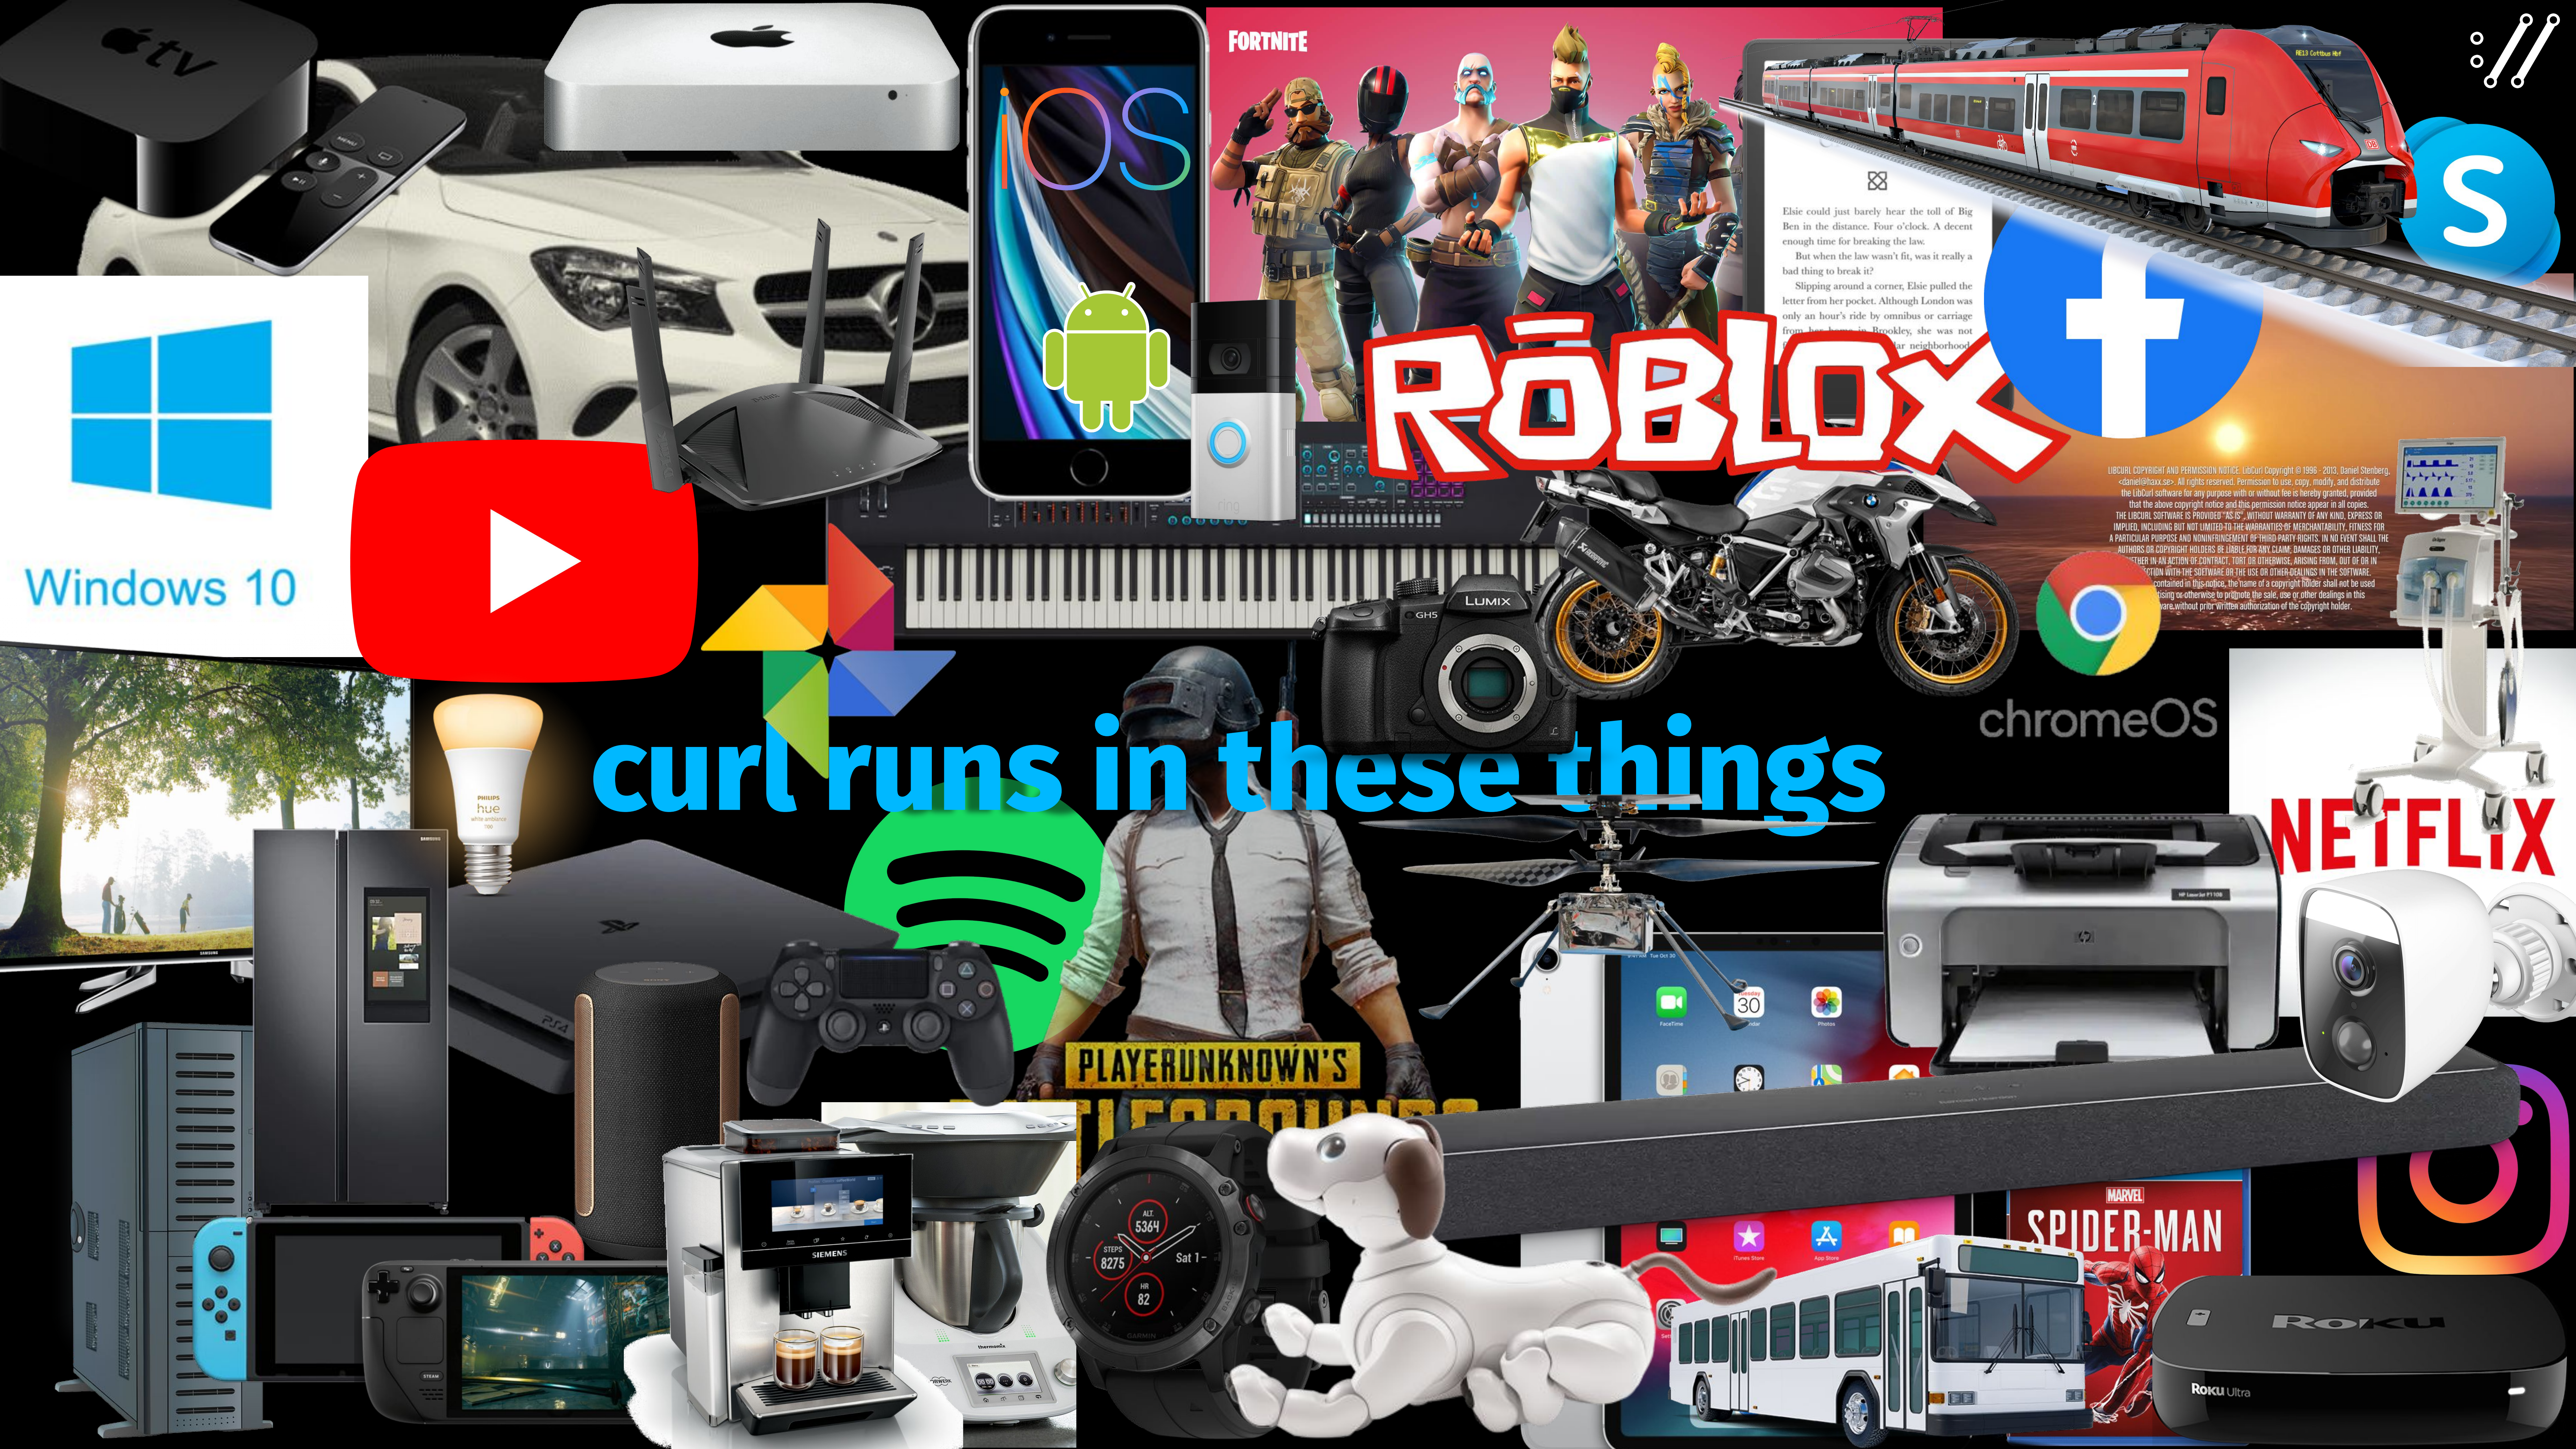

Elsie could just barely hear the toll of Big Ben in the distance. Four o'clock. A decent enough time for breaking the law.

But when the law wasn't fit, was it really a bad thing to break it?

Slipping around a corner, Elsie pulled the letter from her pocket. Although London was only an hour's ride by omnibus or carriage from her home in Brookley, she was not f lar neighborhood

LIBCURL COPYRIGHT AND PERMISSION NOTICE. LibCurl Copyright © 1996 - 2013, Daniel Stenberg, <daniel@haxx.se>. All rights reserved. Permission to use, copy, modify, and distribute the LibCurl software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies. THE LIBCURL SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE. contained in this notice, the name of a copyright holder shall not be used tising or otherwise to promote the sale, use or other dealings in this ware without prior written authorization of the copyright holder.

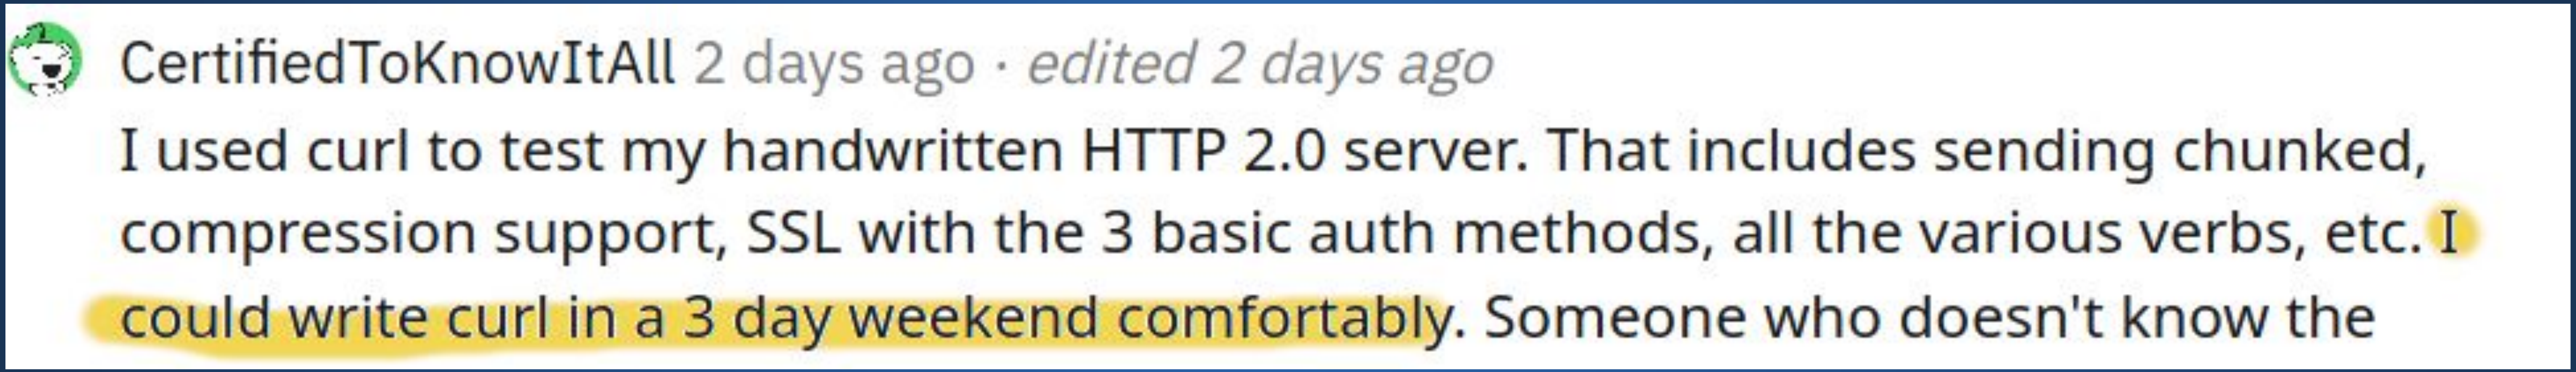

CertifiedToKnowItAll 2 days ago · *edited 2 days ago*

I used curl to test my handwritten HTTP 2.0 server. That includes sending chunked, compression support, SSL with the 3 basic auth methods, all the various verbs, etc. I could write curl in a 3 day weekend comfortably. Someone who doesn't know the

Daniel Stenberg

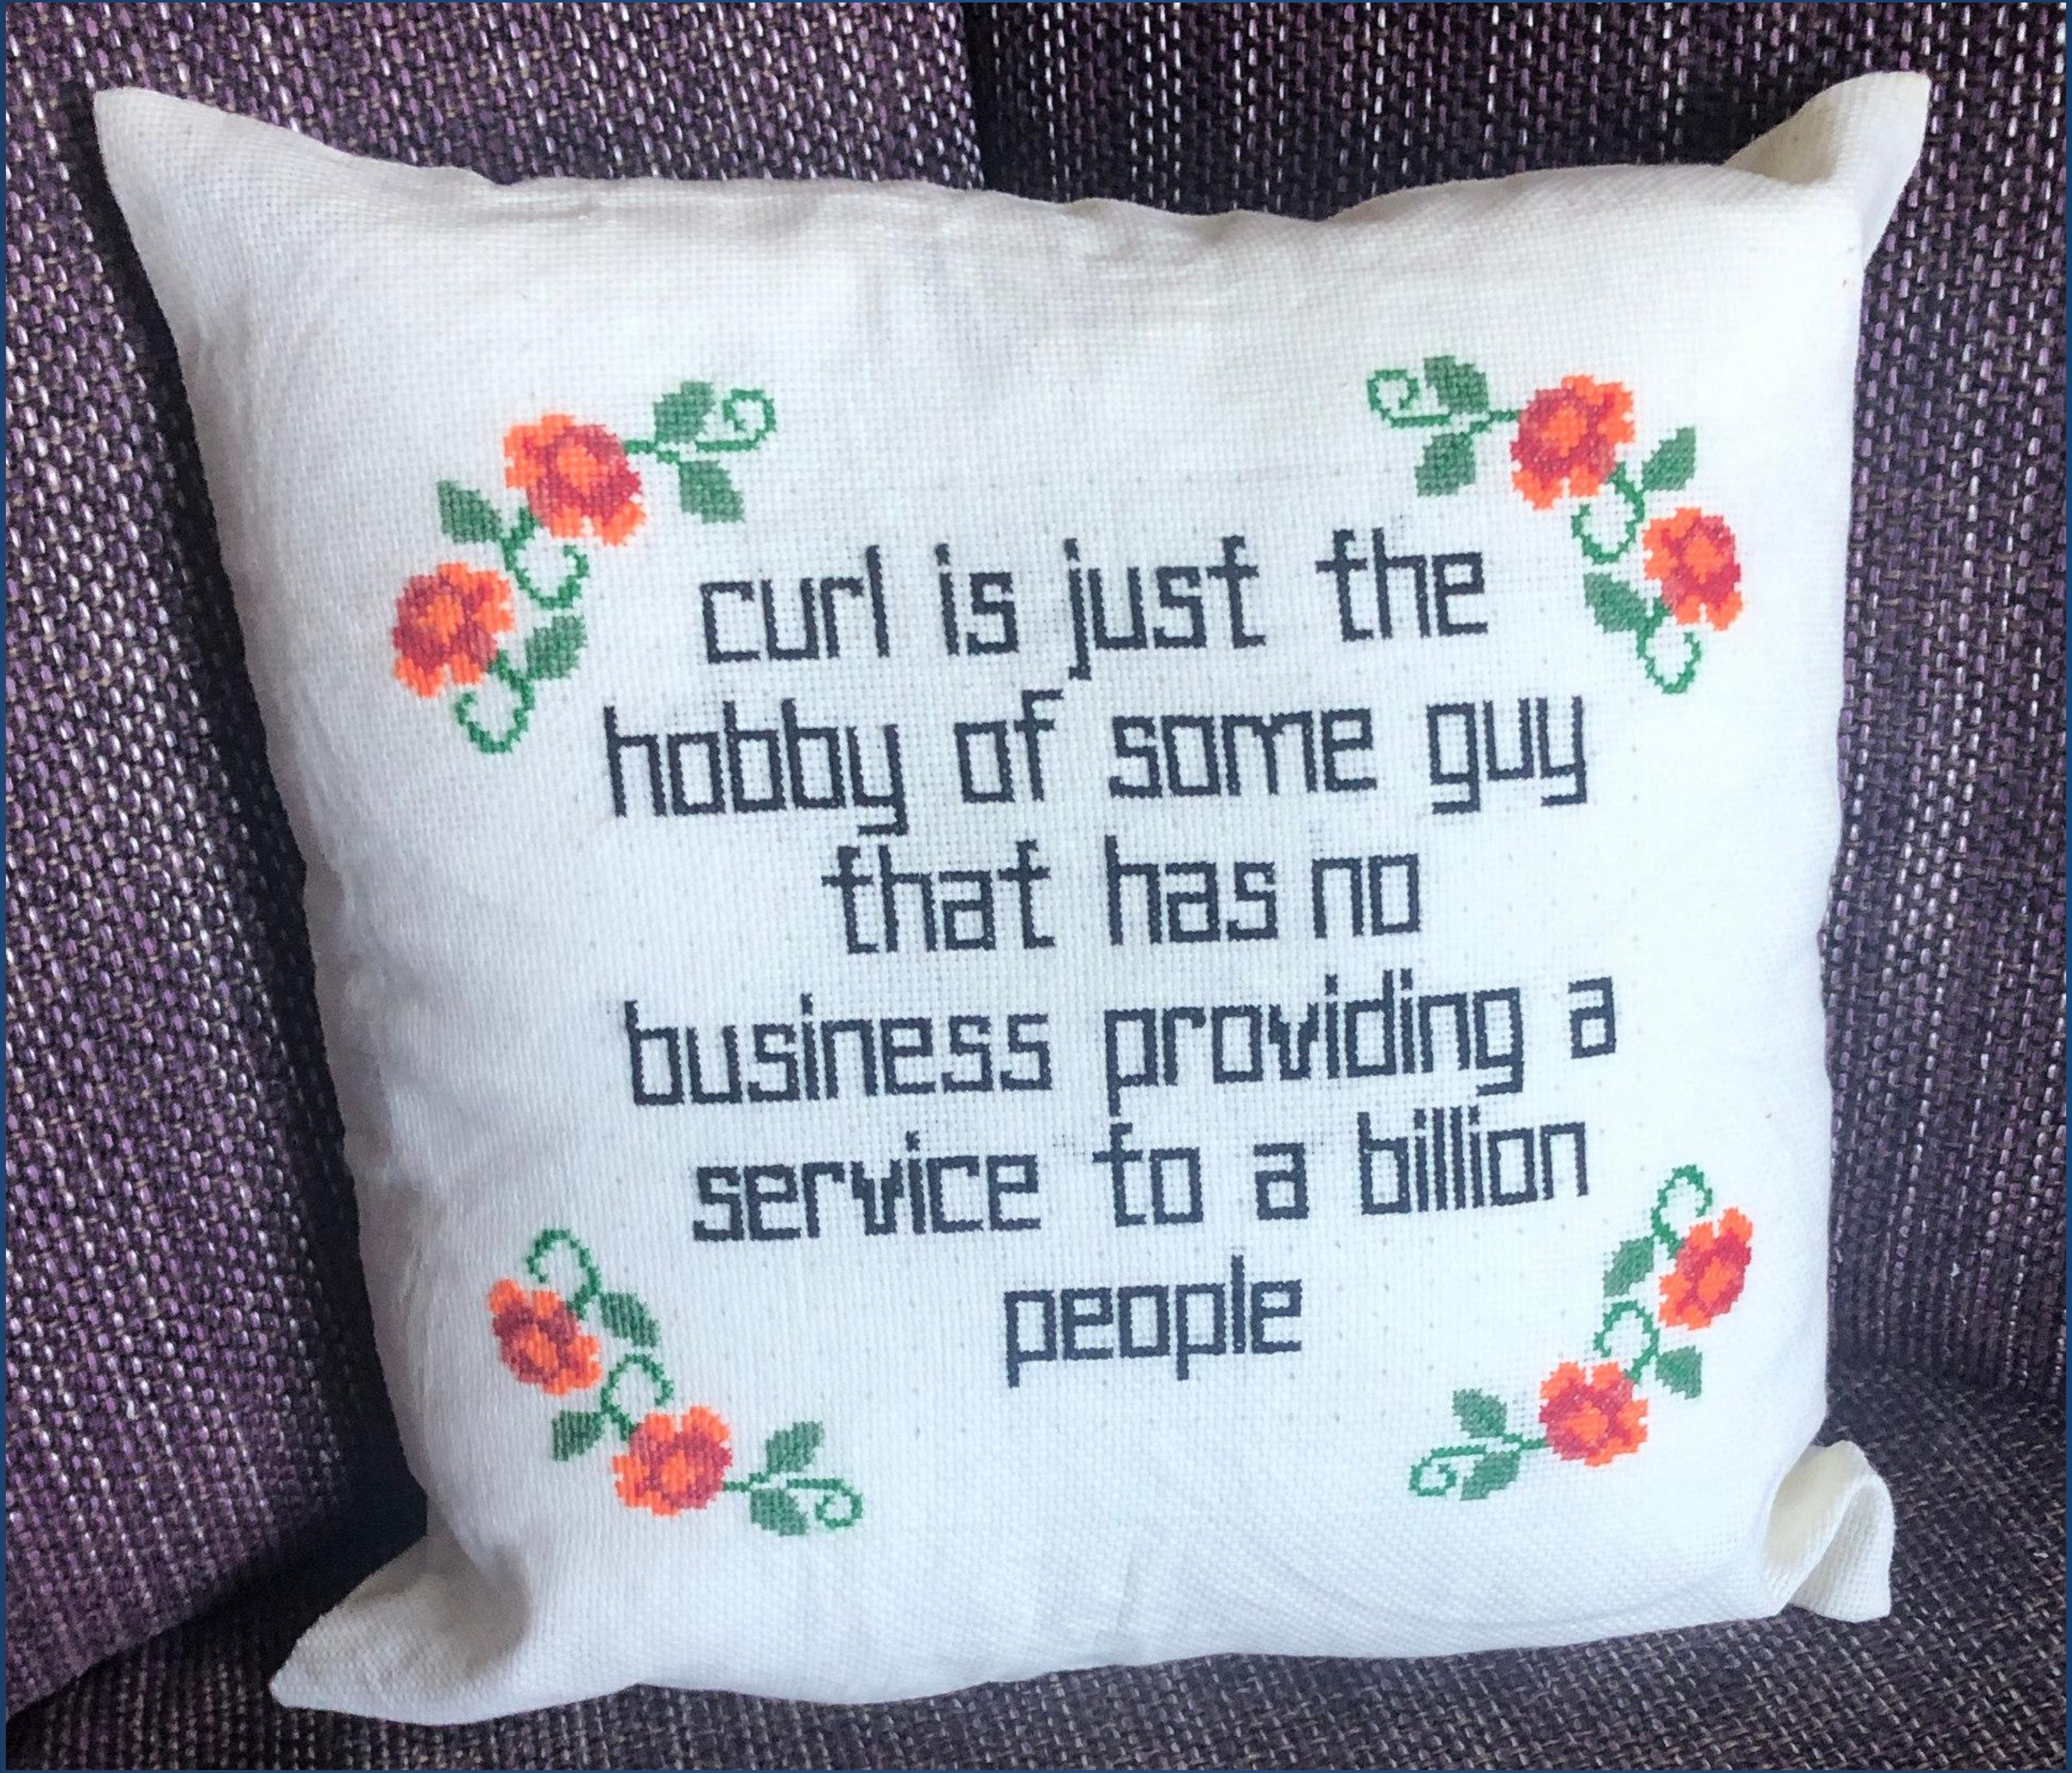

Modern digital infrastructure is to a large degree built on layers and layers of Open Source

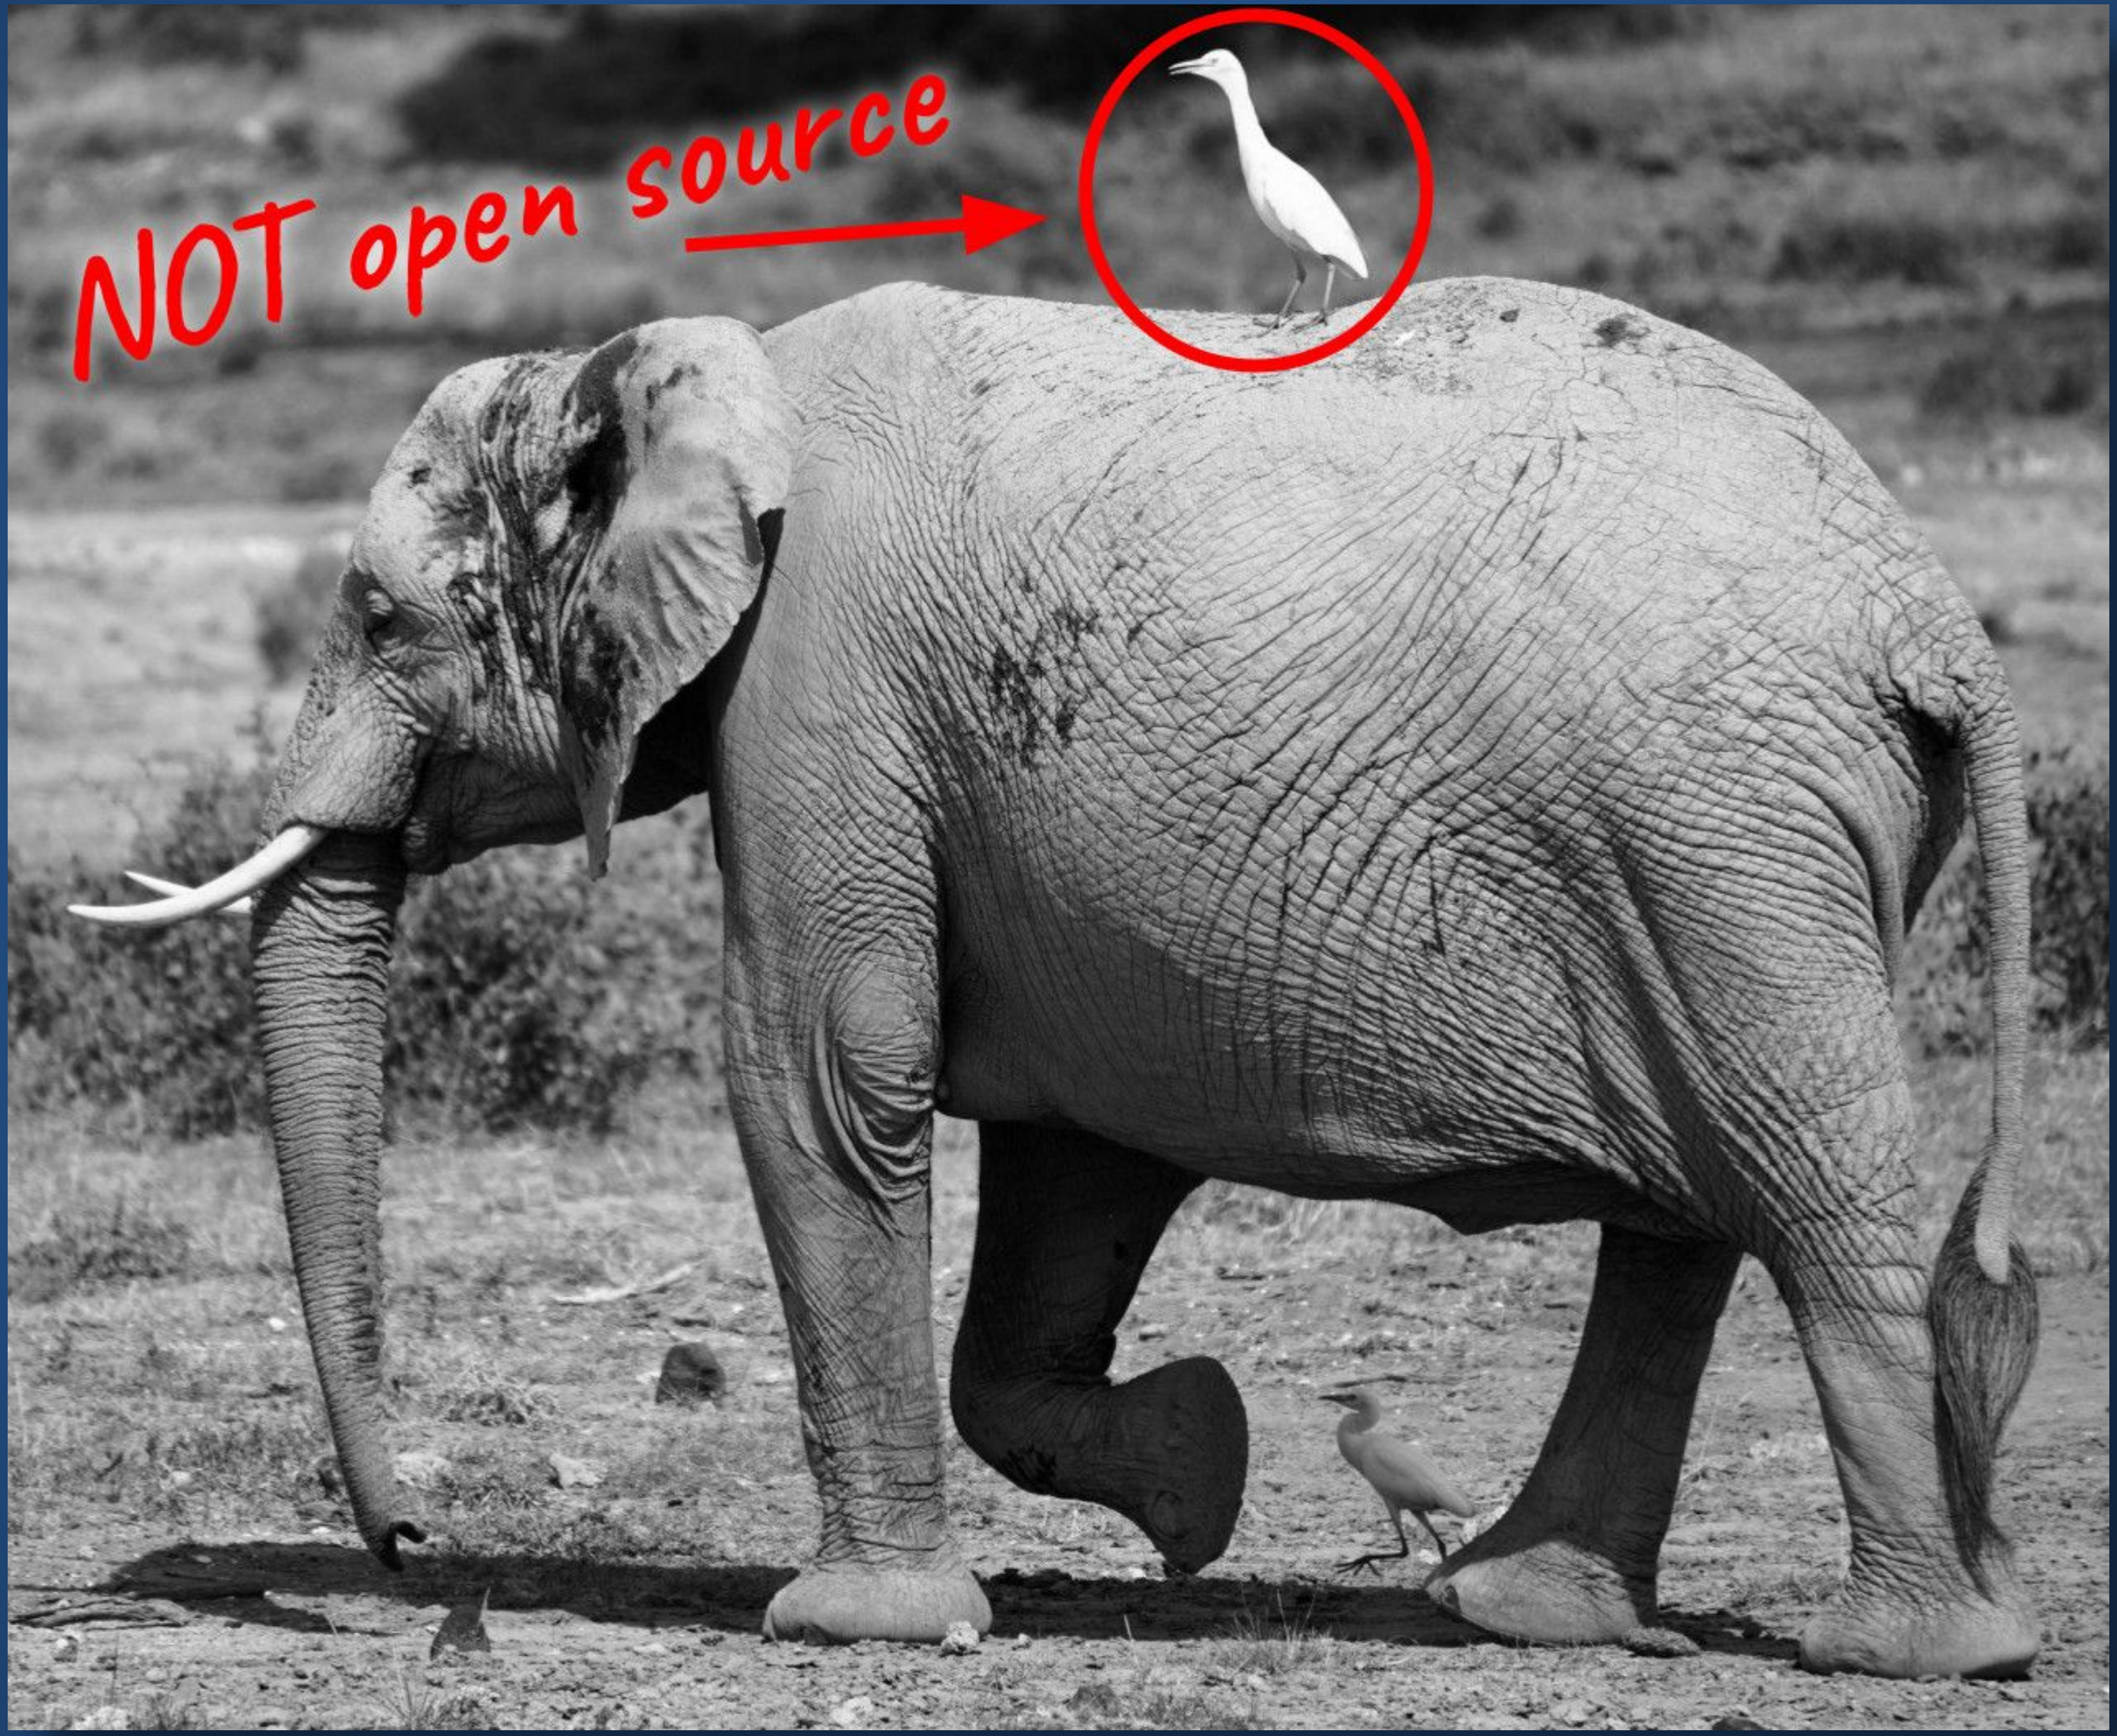
NOT open source

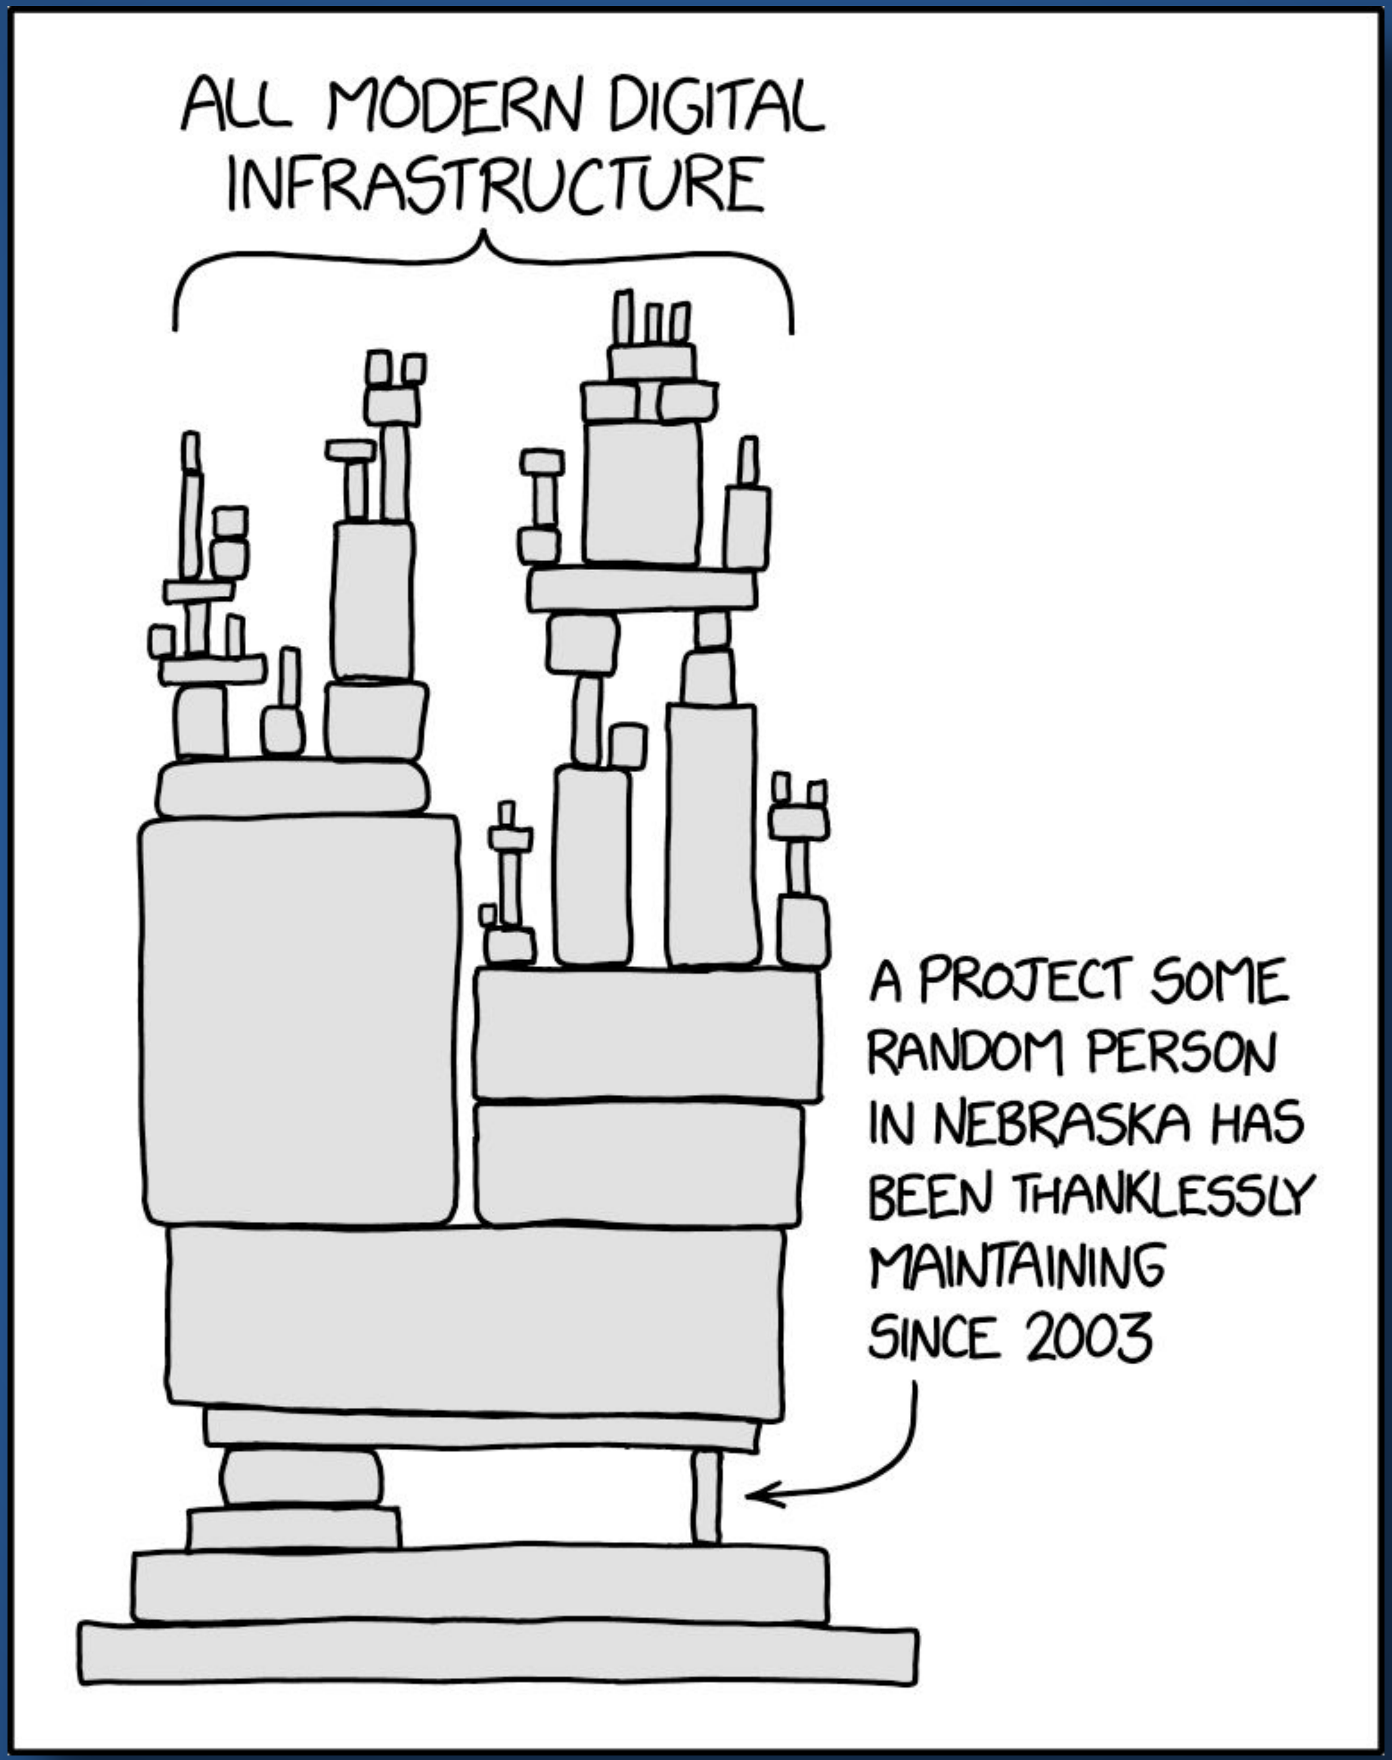
ALL MODERN DIGITAL INFRASTRUCTURE
A PROJECT SOME RANDOM PERSON IN NEBRASKA HAS BEEN THANKLESSLY MAINTAINING SINCE 2003

When digital infrastructure relies on your code, security becomes top priority

Security issues in the project trump all other activities and demand immediate attention

Maintaining Open Source

Most projects have a single maintainer

Many projects are run primarily as a spare time hobby

Many projects are underfunded

Most projects have outstanding tasks

Many maintainers struggle with burnout

an AI slop tsumani

Daniel Stenberg

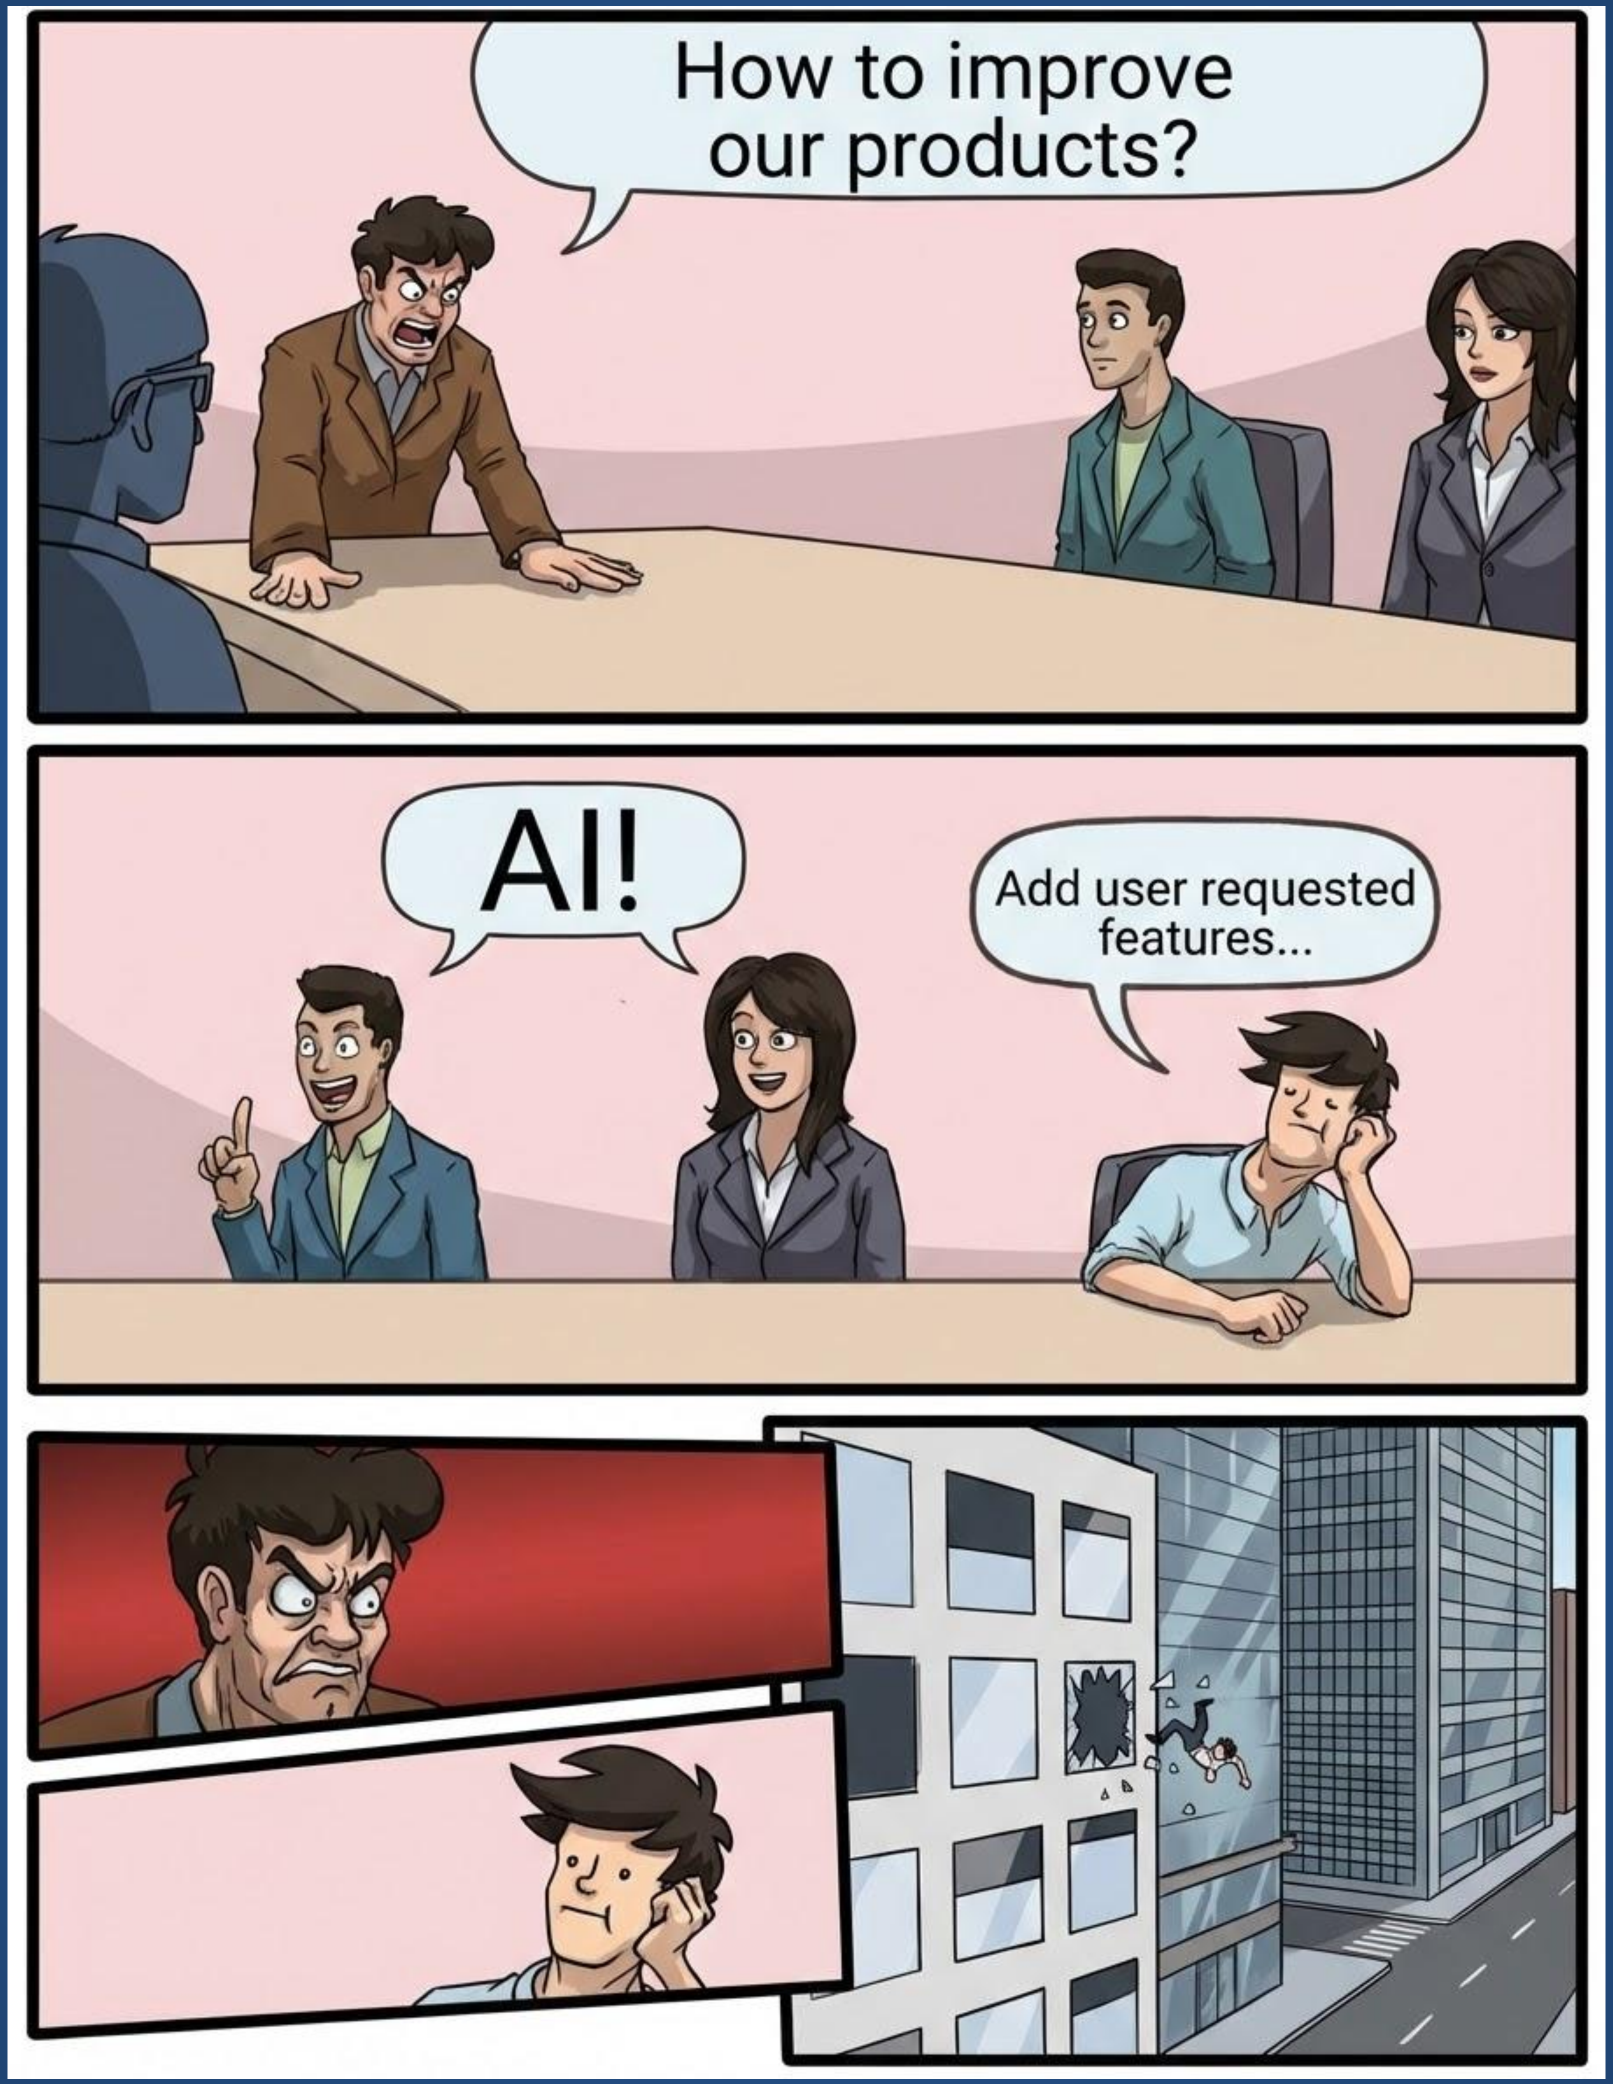

Super easy to ask an LLM to find a security problem to report

At almost zero cost

But is it real?

Does it make us think less?

Daniel Stenberg

Death by a thousand slops

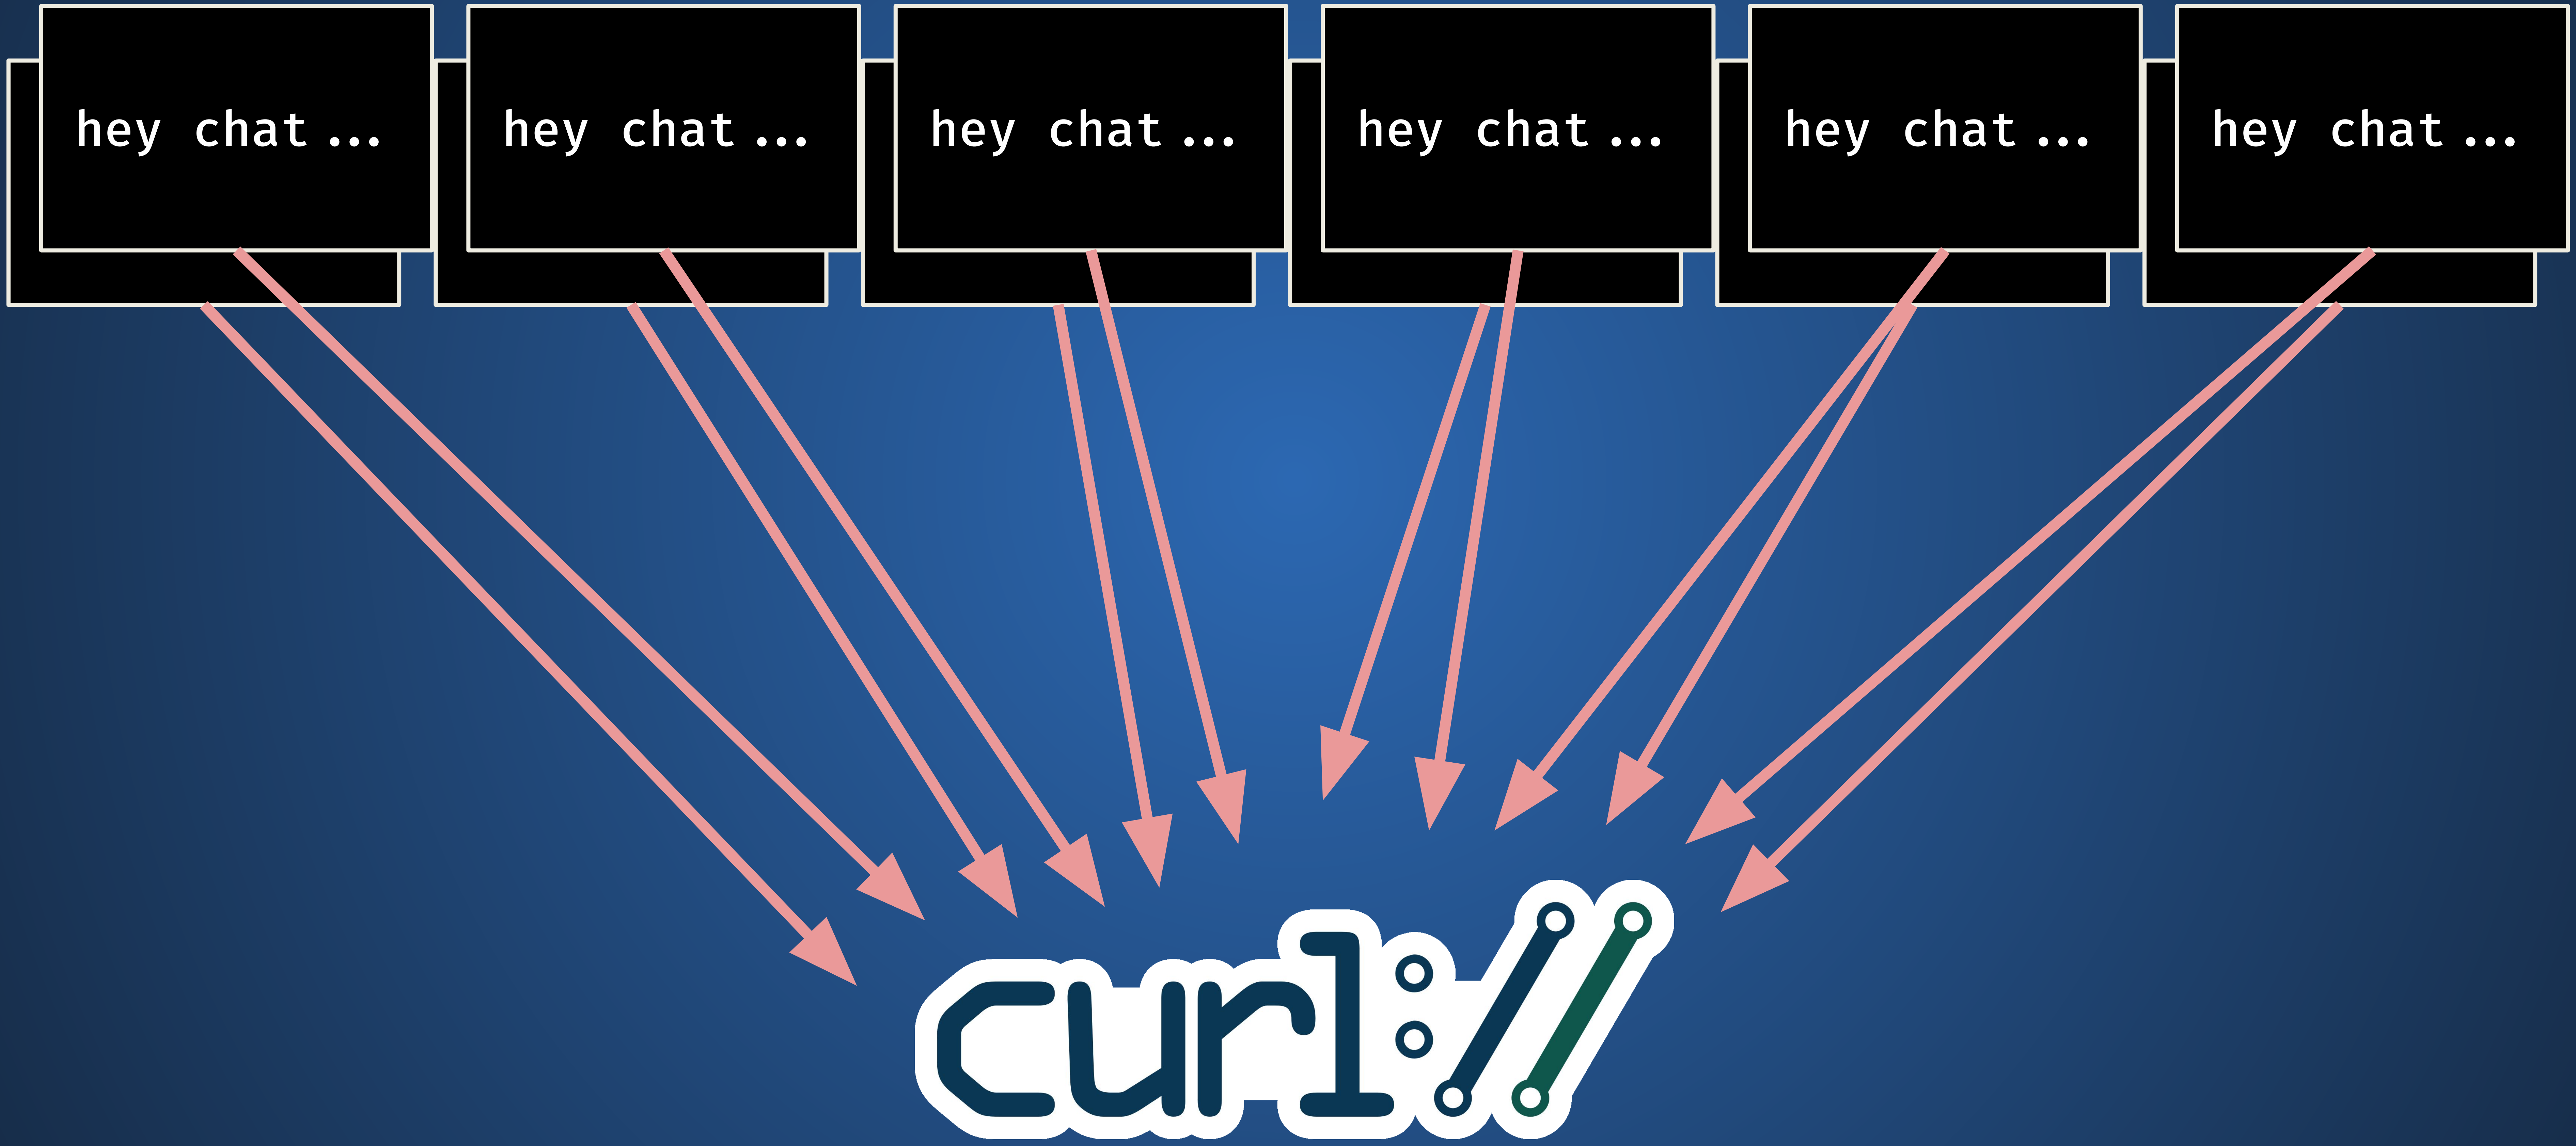

hey chat ...

hey chat ...

hey chat ...

hey chat ...

hey chat ...

hey chat ...

Daniel Stenberg

AI language

too polite

perfect English without typos

Mixed Case To Make It Seem Serious

mdash use

emoji happy

too long - already at first shot

bullet point bonanza

When asked a follow-up question

overly polite and friendly

apologizes a lot

easily loses track and takes off in another direction

replies are also too long

The human involved is just a copy-and-paste proxy

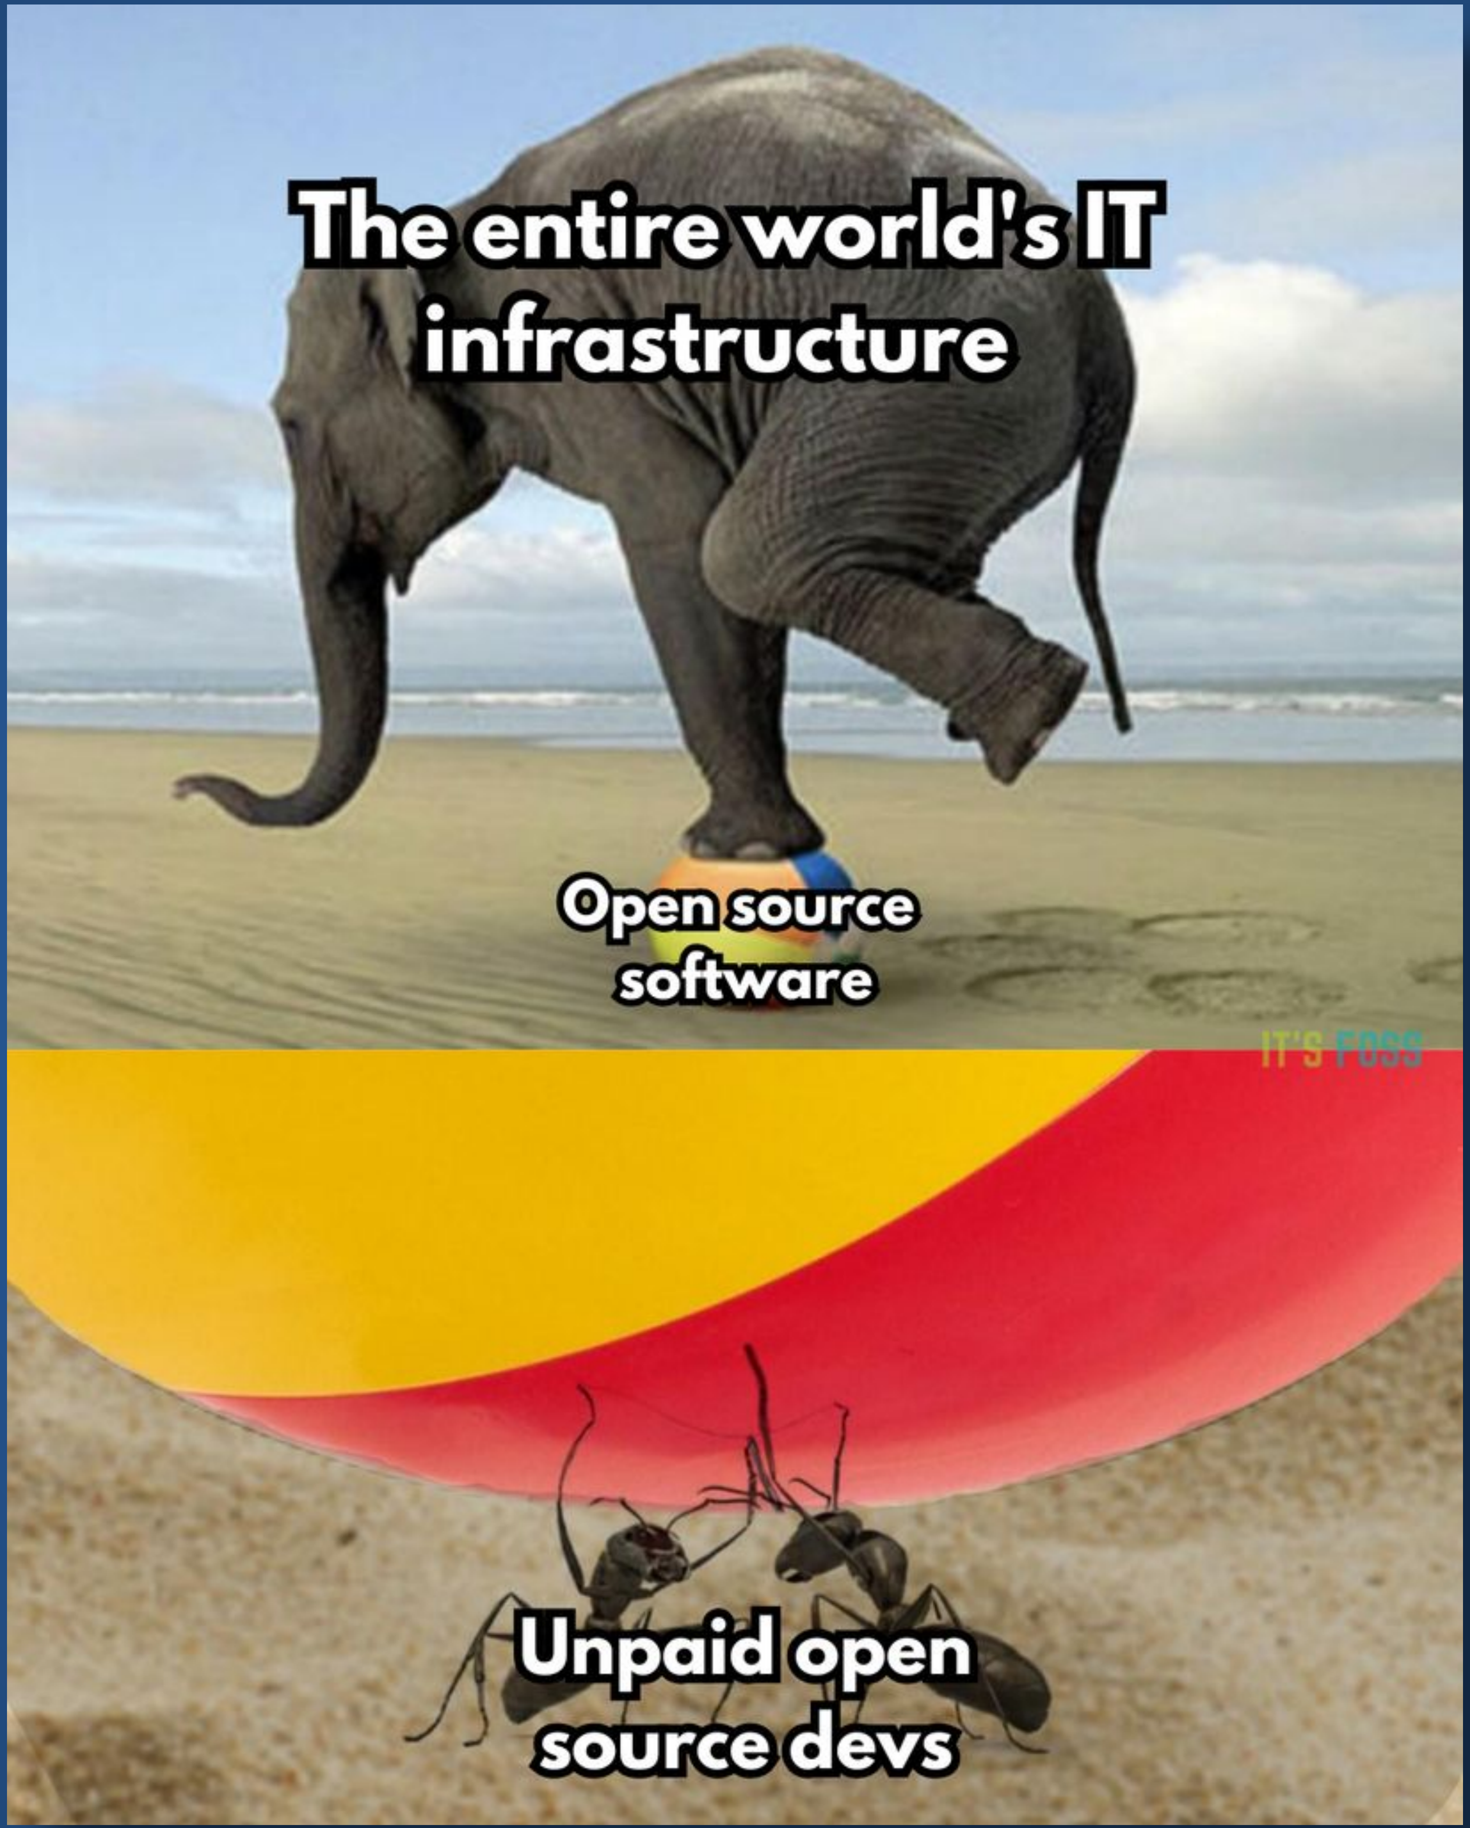
The entire world's IT infrastructure
Open source software
IT'S FOSS
Unpaid open source devs

An AI example

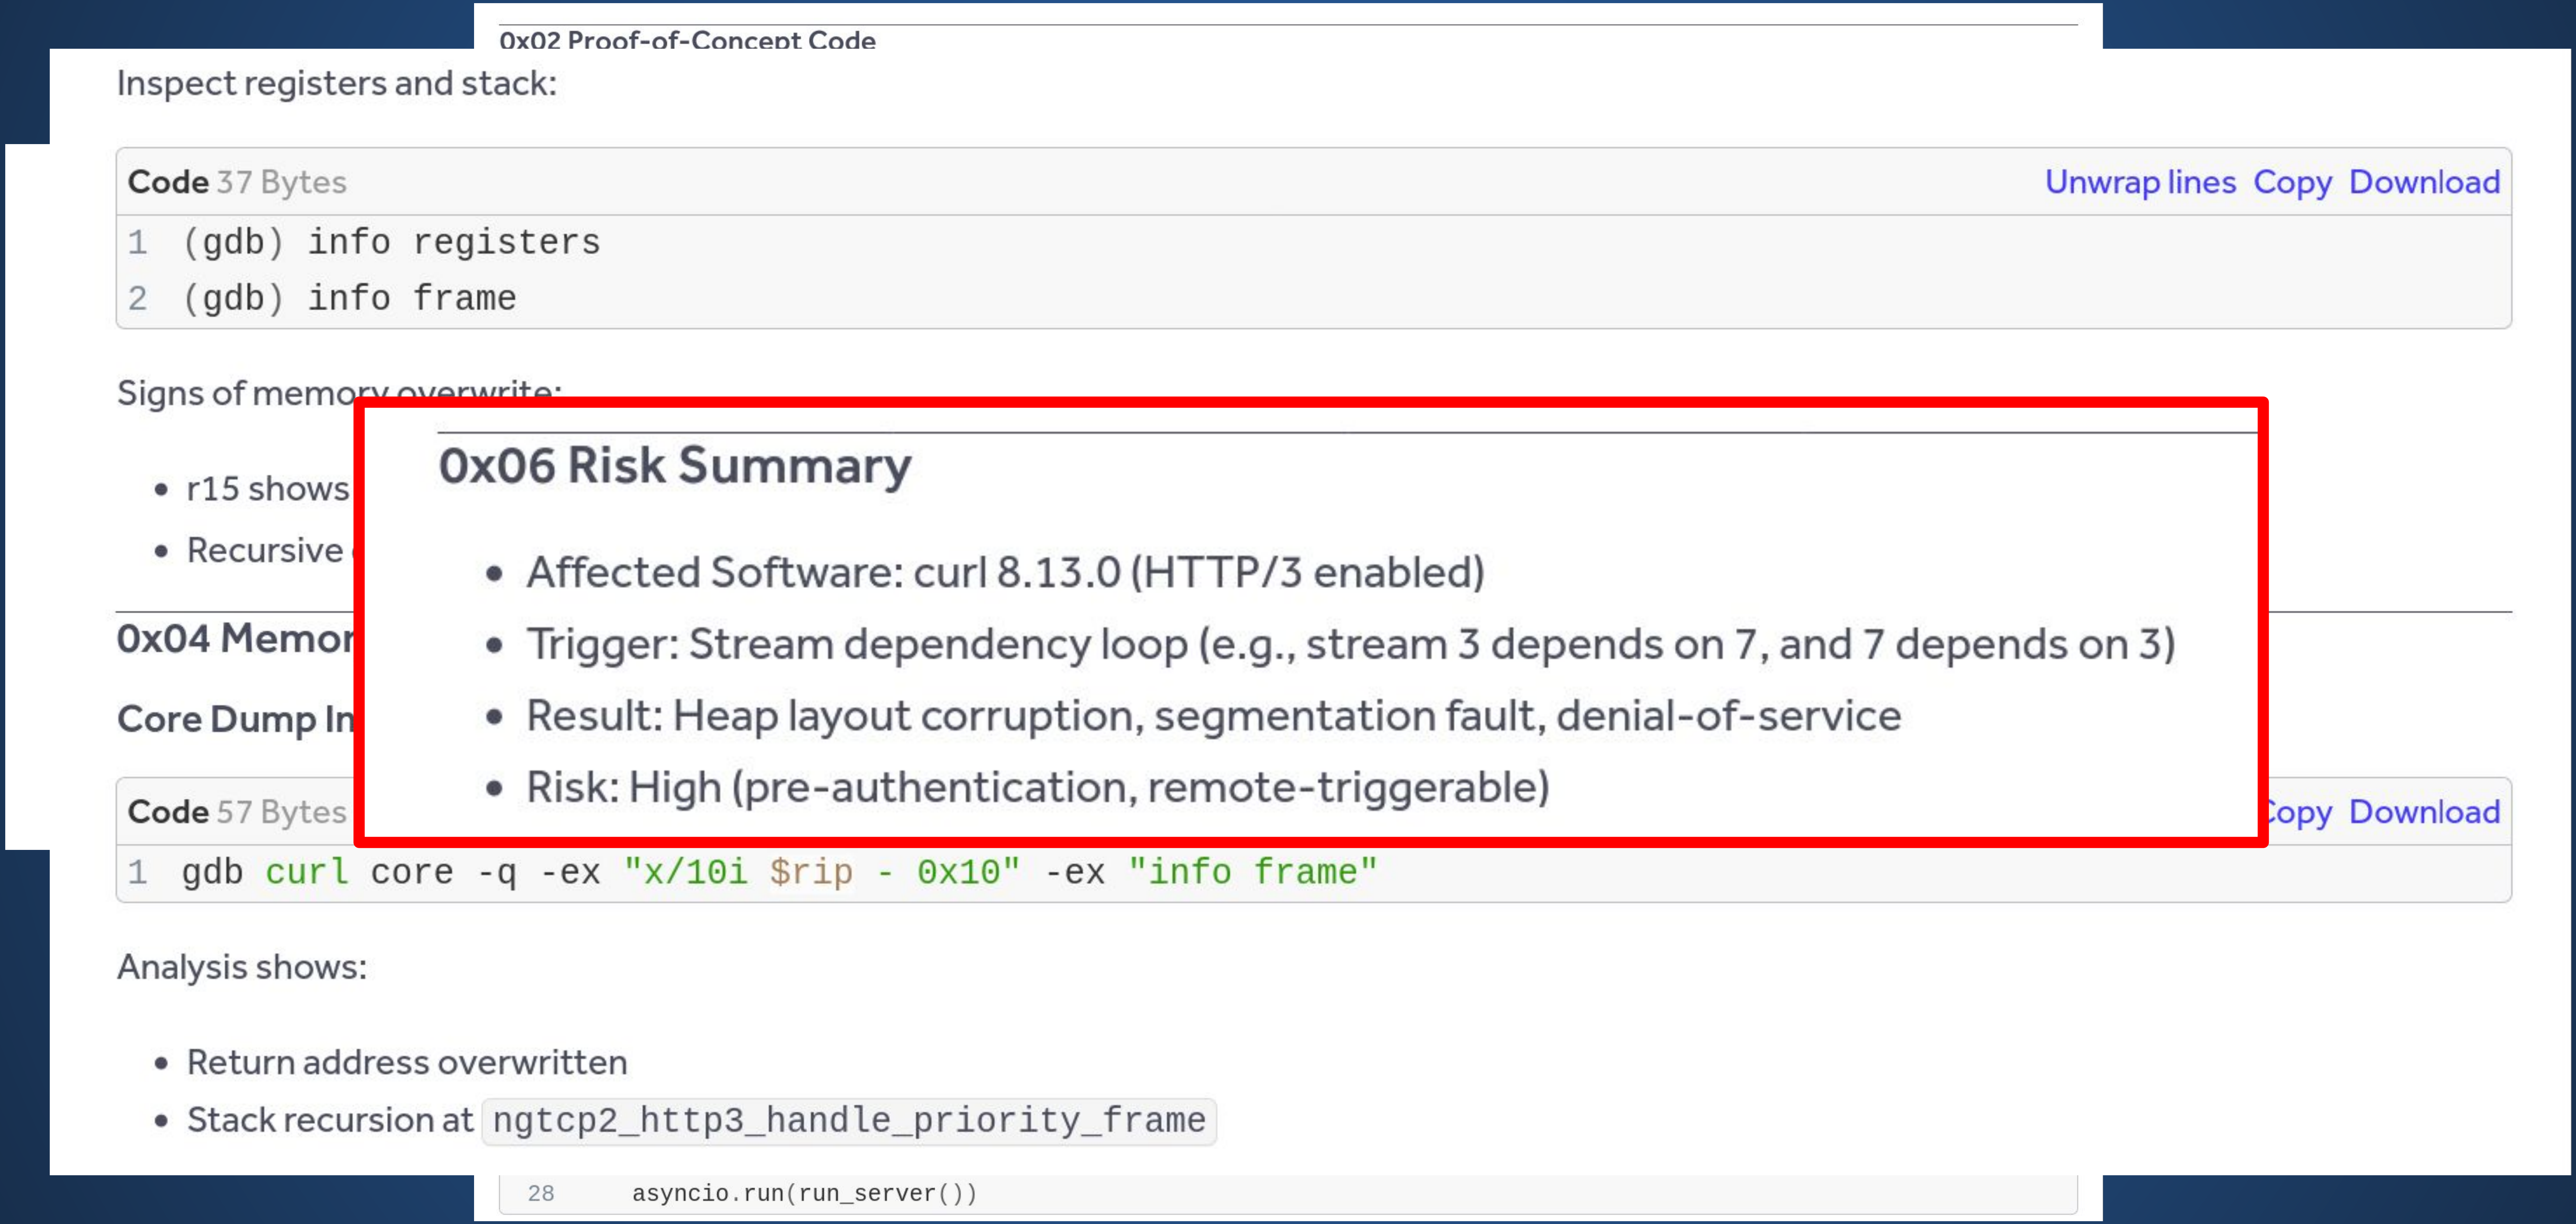

0x02 Proof-of-Concept Code

Inspect registers and stack:

Code 37 Bytes

```
(gdb) info registers
(gdb) info frame
```

Signs of memory overwrite:

- r15 shows
- Recursive

0x04 Memor

Core Dump In

0x06 Risk Summary

- Affected Software: curl 8.13.0 (HTTP/3 enabled)
- Trigger: Stream dependency loop (e.g., stream 3 depends on 7, and 7 depends on 3)
- Result: Heap layout corruption, segmentation fault, denial-of-service
- Risk: High (pre-authentication, remote-triggerable)

Code 57 Bytes

```
gdb curl core -q -ex "x/10i $rip - 0x10" -ex "info frame"
```

Analysis shows:

- Return address overwritten
- Stack recursion at `ngtcp2_http3_handle_priority_frame`

```
28      asyncio.run(run_server())
```

LIES

The mentioned bad function does not exist

The gdb session is made up

The crash does not happen

The shown register contents are made up

Nothing in this report is real

terror reporting

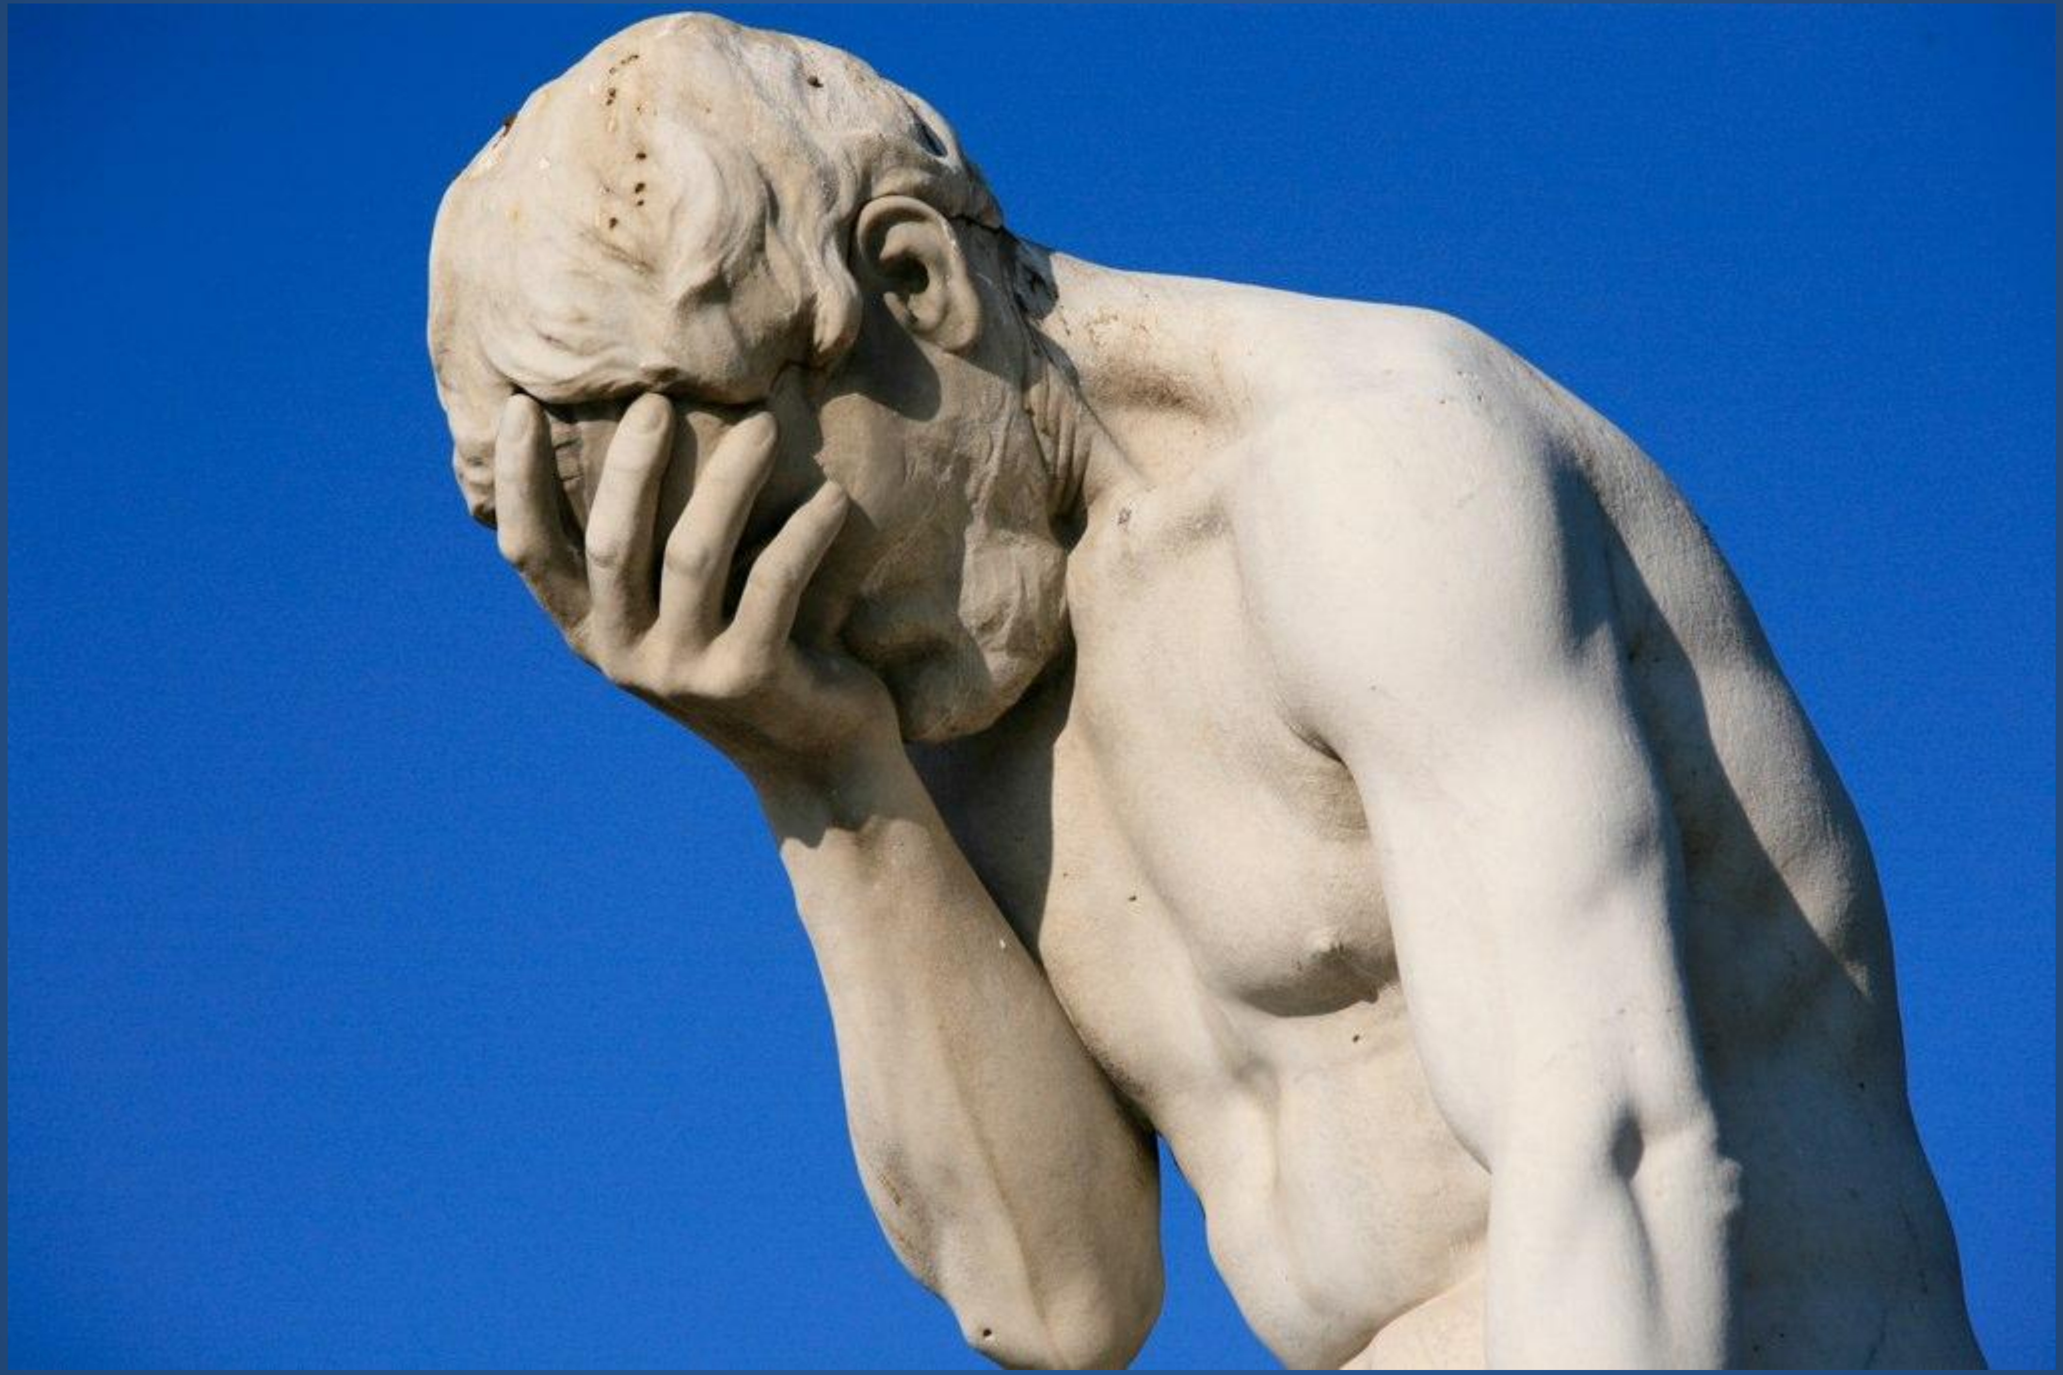

A total waste of time and energy

Stupidity is not AI exclusive

“git repository found”

“information disclosure for… “ (something in the git repo)

“arbitrary file read via file://”

this tool told me [this] is a problem

What does the DDoS attack lead to?

Reduced activity elsewhere

Risks us rejecting real reports

Impacts our sleep, life and mental health

May impact project quality

Which can become a supply chain issue

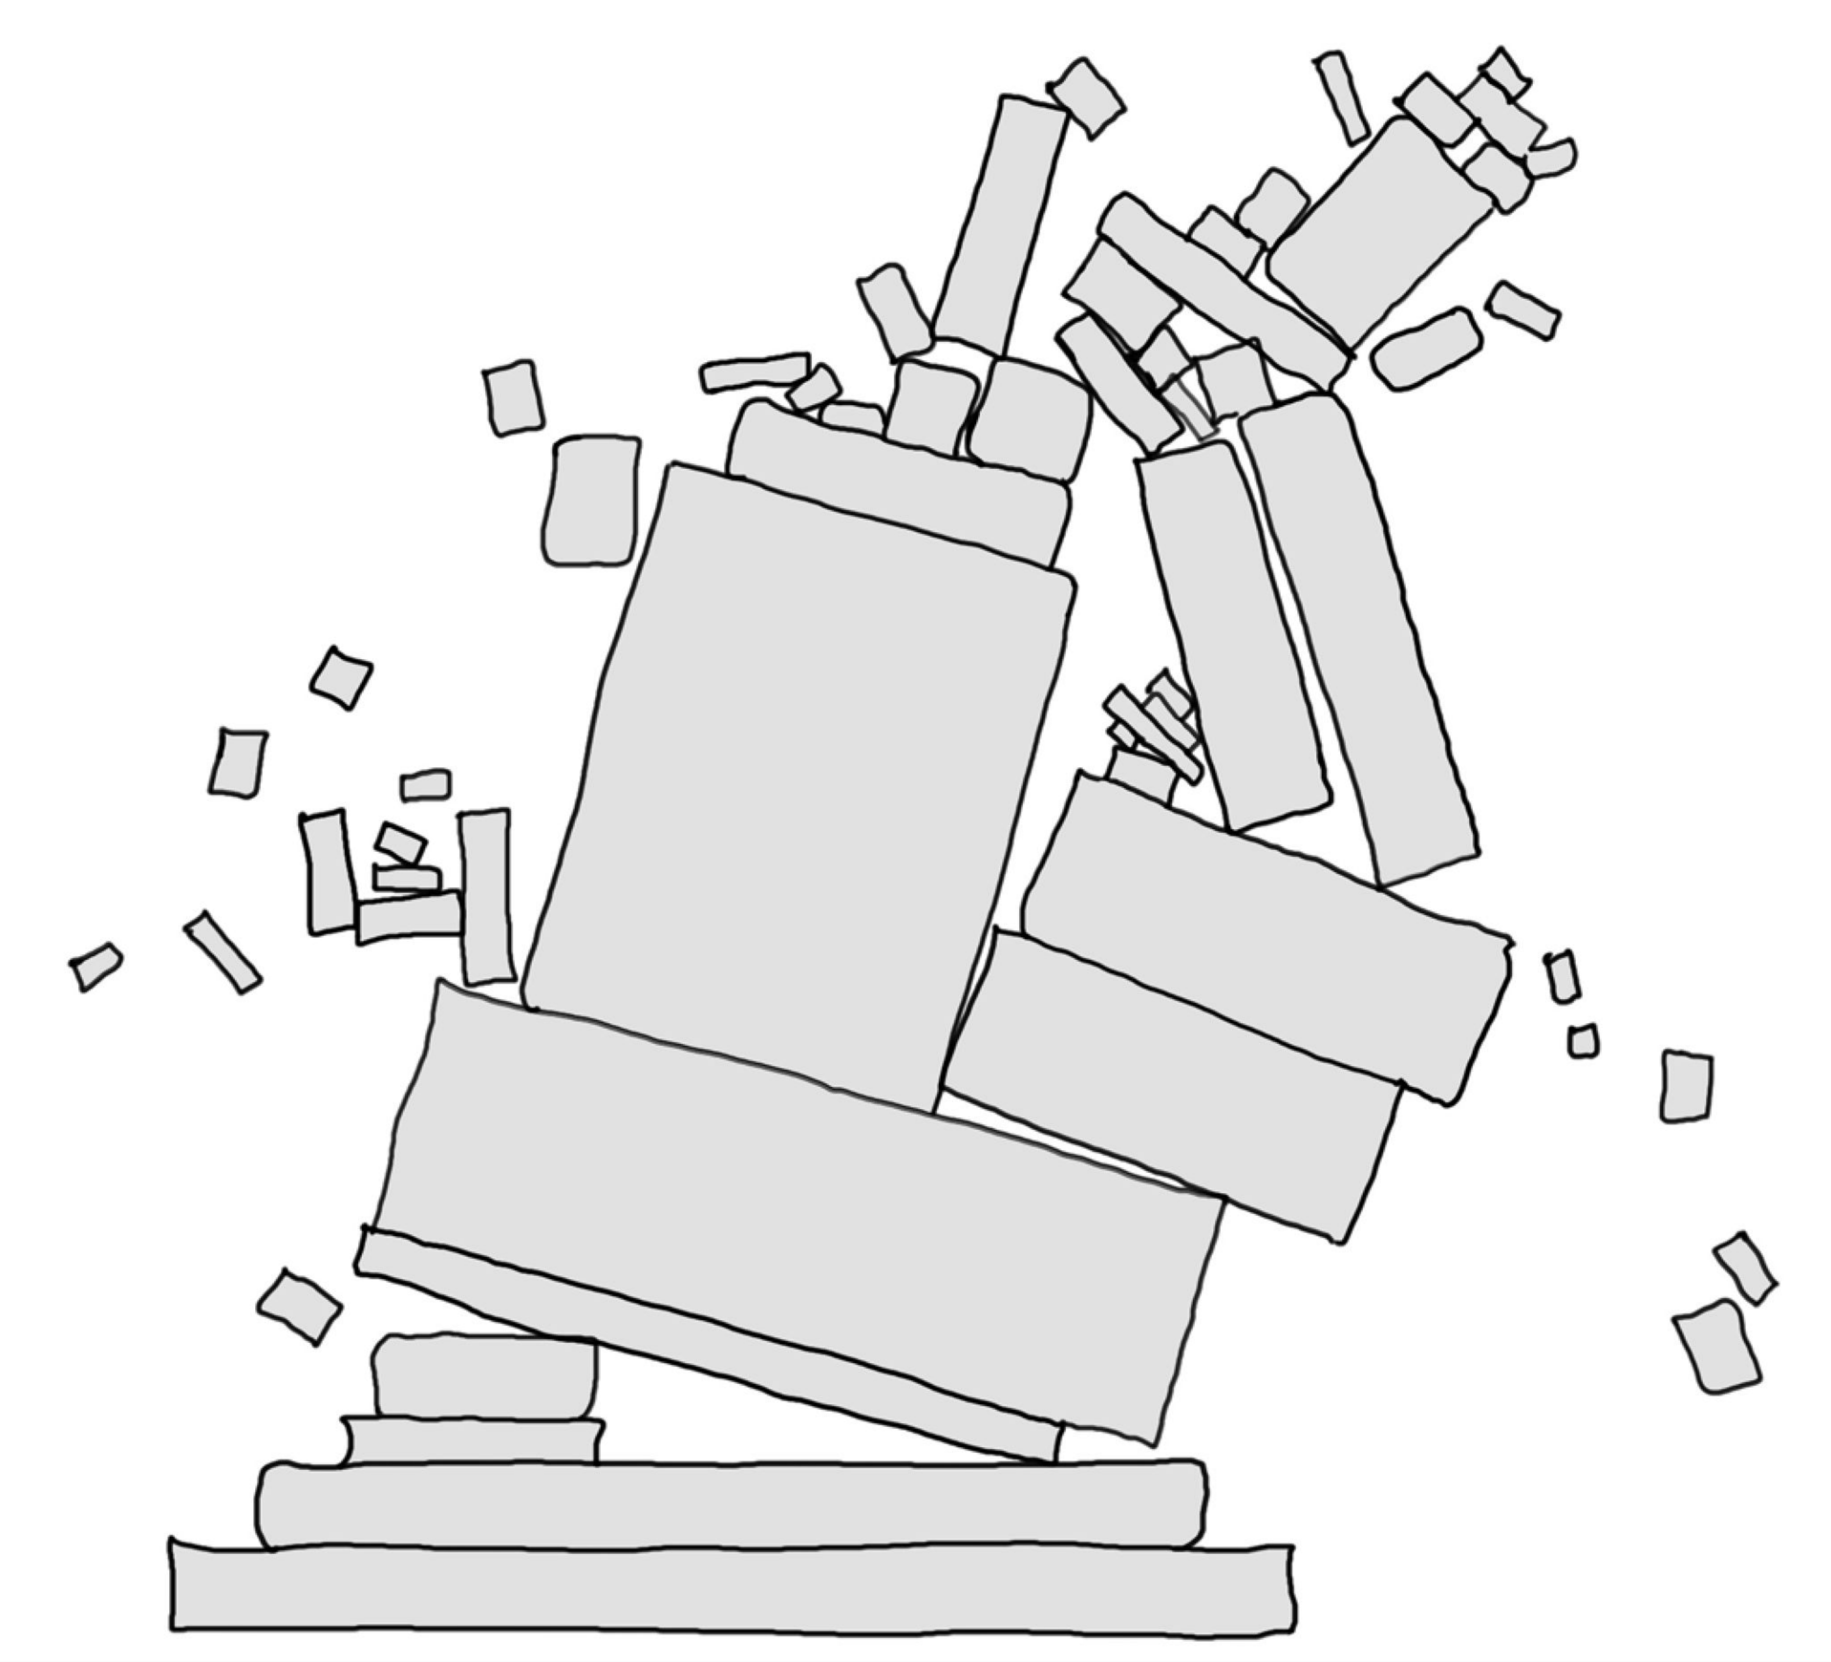

Picture by Adam Leventhal

Why?

The money

Many reporters think they actually have a case

People believe AIs can actually do this

How often does it happen?

Too often

30-70% of submissions

Exact rate is hard to assess

Did not happen two years ago

What we do

Ban the reporter instantly

Require upfront notification about AI use

We want to remain approachable and open

We want everyone to be able to report issues

Public shaming

Abuse

Not really about AI, but the abuse

AI makes it easy and cheap

AI marketing mislead people

Users cannot discern AI lies from truths

Human created slop is also a problem

The curl bug-bounty is now officially shut down

AI analyzer tools

For the good or for the bad

We work with Aisle Research, ZeroPath and more

Easily lead to issue overload

They find many mistakes no other tools find

Might be (ab)used by adversaries

Still need a human brain to filter, assess, fix

In curl we have fixed 100+ issues found by AI tools

AI tool findings

“this code does not comply with the protocol spec”

“this comment says something the code disagrees with”

“this implementation forgot an edge case”

“there is an edge case for this that is not tested”

code that “normal” analyzers cannot scan

takes 3rd party libs into account like no “normal” analyzer

AI for code review

AI powered code review

We toy with GitHub Copilot, Augment Code and Aisle Research bot

Provide a first review fast

Find mistakes pretty well

Easy to dismiss when you disagree like with human review

The AI tools find different things to remark on

AI powered code review

Humans are **not very good** at review

Easy to use **when cheap**, but then?

Proper **test coverage** is still **more important**

Compare test failures the review **did not detect**

AI writing code

AI code producers

Not impressed by AI powered editors, copilot or zeropath

Having an eager junior throwing ideas at you might help

Not used (officially) in or by the curl project

AI scraper overload

curl.se bandwidth spend

fastly®

serves 75 terabytes/month

over 4000 requests/second

tarball downloads are < 0.01% of the requests

Future

AIs will improve

AI companies will continue selling the myths

Humans will most likely not improve

AI augments human activities *in all directions*

Maybe they soon start to charge us what it costs https://daniel.haxx.se/ai-slop

Warning: mind-numbingly stupid

QUESTIONS?

License