

Web Security is Too Hard

It started innocently enough. I saw a tweet about a new product offering from one of my favorite companies, Cloudflare.



Michelle Zatlyn ✓

@zatlyn



I just reserved my Cloudflare Wallet tag: zatlyn.cloudflare.pay.

Reserve yours now at cloudflare.pay

10:06 AM · Aug 4, 2026 · **5,374** Views



4



3




34



1




Neat! I clicked through to the site and there it is:




And *huzzah!*, my preferred handle, @ericlaw is still available. I'd better hurry to claim it before someone else gets it!

Since I'm already a long-time Cloudflare user, I just need to sign in. That makes sense, how else will they bind the handle to my account?



Easy peasy. I'm in. Looks like there's just one more step, I gotta authorize the new feature?



But wait a sec!

This looks exactly like one of those [Consent Phishing attacks](#) that have been so popular over the last few years!

And wait, why is the entry point on `cloudflare.pay`, a site that *doesn't* already have my credentials, rather than something within the `cloudflare.com` domain which does (e.g. `cloudflare.com/pay`)? There is no inherent technical relationship between a `.com` domain and a `.pay` domain. Domain names under the `.pay` sTLD are available to [anyone with \\$20](#) (unlike, e.g. [.bank](#) which requires more vetting), so there's nothing that would stop me from registering my own `cloudflarepayments.pay` domain name in just a few minutes.

And why doesn't Cloudflare's permission site recognize its own company's feature? And that green checkmark looks suspicious as heck— an attacker could probably just shove that emoji inside their misleading display name, the same way that folks trying to [phish Microsoft email accounts](#) use misleading app names and icons:



**Let this app access your
info?**


Fake Outlook OAuth phishing request

The guys at Cloudflare are geniuses who know their stuff. This has *got* to be an attack. It's a clever one — I was feeling such a [sense of urgency](#) because I wanted to “win” the race to get my desired handle. *Very very clever!*

Unfortunately, the Cloudflare permission page doesn't follow [best practices](#), so there's no “Report suspicious request” link I can use to let the Cloudflare folks know that their customers are under attack.

Let me go back to my Cloudflare dashboard and try to get to the Wallet feature from its sidebar. Hrm. It's not there. Now, Wallet purports to be “a new feature”, so maybe the Dashboard just isn't updated yet. A search of the docs turns up nothing. Let's ask the AI agent in chat.

The very first thing the chat agent wants is access to my account:



This feels a little weird, but the page is still `cloudflare.com` so I guess I can give the thing access to things it already has access to. Weirdly, the AI agent first proposes that I grant it full control rather than read only access, which feels like a failure of the principle of least privilege, but I don't actually need to ask an account specific question anyway. After granting read permission, the agent allows me to ask my question:

Oh, wow. Cloudflare says it really *is* an attack! Let's [report the phish right away!](#)

A few minutes later... womp womp...

Oh dear.

After a few minutes of further frantic searching, it turns out that this is, in fact, a legitimate new Cloudflare product and a legitimate site, despite giving every indication of being a clever phishing attack.

It further turns out that that suspicious green checkmark is not part of the app's untrustworthy display name but instead a (poorly placed) security UI element that a user is expected to hover over to get the security details:

Update Aug 6th: The Cloudflare folks have improved this UI today. It now plainly shows the relationship information *by default* and not directly adjacent to the third-party string:

The Cloudflare folks apparently want security issues reported via HackerOne (*which wouldn't let me log in because the Cloudflare CAPTCHA HackerOne uses seems to be broken...*).

When even legitimate websites sometimes act very very phishy, consider how hard it must be for URL Reputation services like Microsoft SmartScreen and Google SafeBrowsing to block malicious sites without false positives, especially in a world where millions of new sites are added to the web every week.

Lessons

Web Developers, please follow every best practice, I'm *begging* you:

- Host apps and content under your trusted domain name (e.g. `cloudflare.com/pay` or `pay.cloudflare.com`. *If you must add a new name, link*

to it directly from a page on your trusted domain name.

- Show relevant security information in a [trustworthy place](#) when asking the user to make security decisions.
- Make it [trivial to report scams](#), in context (e.g. on the permission request page).
- Ensure permission screens follow [Best Practices to Avoid GestureJacking](#) to prevent users from unknowingly granting permissions.
- Test your security reporting flows to ensure they are monitored and function correctly.

Users: Try to stay safe out there. Think before you click, and if all else fails, [wait](#).

Security Geeks: Never blame the victim— they’ve got an impossible job.

-Eric

 [ericlaw](#)  [2026-08-04](#)  [browsers](#), [security](#), [storytelling](#), [web](#)
 [best-practices](#), [browsers](#), [phishing](#), [SmartScreen](#)

Published by ericlaw

Impatient optimist. Dad. Author/speaker. Created Fiddler & SlickRun. PM @ Microsoft 2001-2012, and 2018-, working on Office, IE, and Edge. Now working on Microsoft Defender. My words are my own, I do not speak for any other entity. [View more posts](#)

2 thoughts on “Web Security is Too Hard”

midp

[2026-08-06 at 01:37](#)

Looks like they already fixed their UI elements, likely right after reading this post.
Props to you for your due dilligence though :-)

★ Loading...

[Reply](#)

[HisVirusness](#)

2026-08-06 at 04:42

They could have easily just used the “.pay” domain to forward to their trusted domain. That’s such a planetarian-level oversight.

★ Loading...

[Reply](#)

Leave a Reply
