19 AUGUST 2018 / TECHNICAL

Finding The Real Origin IPs Hiding Behind CloudFlare or TOR

Hidden services and the effectiveness of CloudFlare or any similar service live from hiding the origin servers IP. Simple small mistakes can reveal the IP. This allows attacking a website that uses CloudFlare directly (bypassing the WAF, Rate Limits, DDoS Protection and much more) or even un-hiding a TOR hidden services operator identity. The mistakes depend on what type of service or technology you are working with, not all methods work for every technology (e.g. MX records don't exist for TOR hidden services).

1. SSL certificates

1.1 Using a given domain name

You are hosting a controversial service on xyz123boot.com. Your origin server IP is 136.23.63.44. CloudFlare is providing you with a DDoS Protection, Web Application Firewall and a couple of other services, that protect your project from the people, that would like to see your project offline. Your web server supports SSL and has a certificate, so the communication between CloudFlare and your server is encrypted just like the communication between your users and CloudFlare (i.e. no Flexible SSL). Everything else would be a false sense of security. So far so good.

The problem is the fact you're also **exposing the SSL certificate when directly connecting to your IP** on port 443 (https://136.23.63.44:443). Scanning 0.0.0.0/0, the whole internet, on port 443 for a certificate valid for xyz123boot.com will give your web servers IP to the attacker.

Censys is doing the scan for you. The only thing you have to do is translating the above search terms described in words into actual search queries.

Certificates for xyz123boot.com: `parsed.names: xyz123boot.com`
Just show valid ones: `tags.raw: trusted`

Combining multiple parameters on Censys can be done by using simple Boolean logic.

Search term: `parsed.names: xyz123boot.com and tags.raw: trusted`

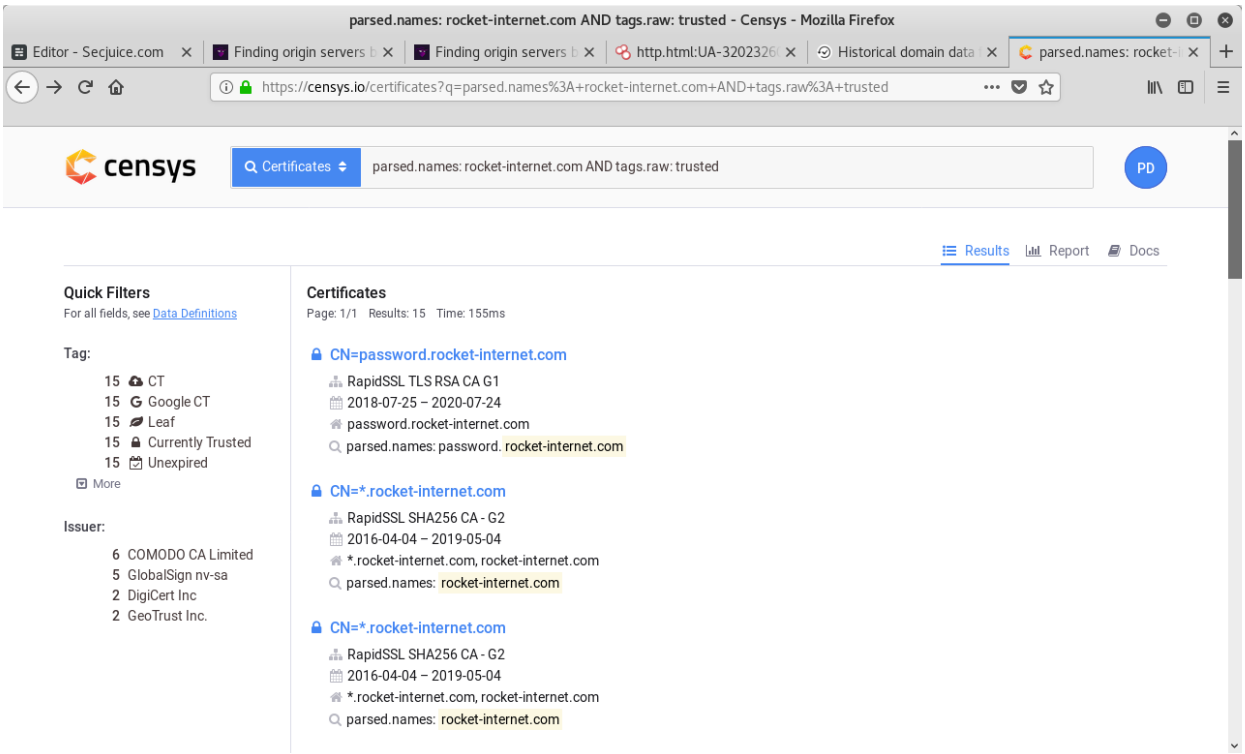

https://censys.io/certificates?q=parsed.names%3A+rocket-internet.com+AND+tags.raw%3A+trusted

parsed.names: rocket-internet.com AND tags.raw: trusted

Quick Filters

For all fields, see Data Definitions

Tag:

15	CT
15	Google CT
15	Leaf
15	Currently Trusted
15	Unexpired

Issuer:

6	COMODO CA Limited
5	GlobalSign nv-sa
2	DigiCert Inc
2	GeoTrust Inc.

Certificates

Page: 1/1 Results: 15 Time: 155ms

CN=password.rocket-internet.com
RapidSSL TLS RSA CA G1
2018-07-25 – 2020-07-24
password.rocket-internet.com
parsed.names: password. rocket-internet.com

CN=*.rocket-internet.com
RapidSSL SHA256 CA - G2
2016-04-04 – 2019-05-04
*.rocket-internet.com, rocket-internet.com
parsed.names: rocket-internet.com

CN=*.rocket-internet.com
RapidSSL SHA256 CA - G2
2016-04-04 – 2019-05-04
*.rocket-internet.com, rocket-internet.com
parsed.names: rocket-internet.com

Censys will show you all the certificates matching the above criteria, which they found in their scans.

Clicking on the search results one by one, you can open a drop-down with several tools by clicking on "Explore" on the right side. Choose `What's using this certificate?` > `IPv4 Hosts`.

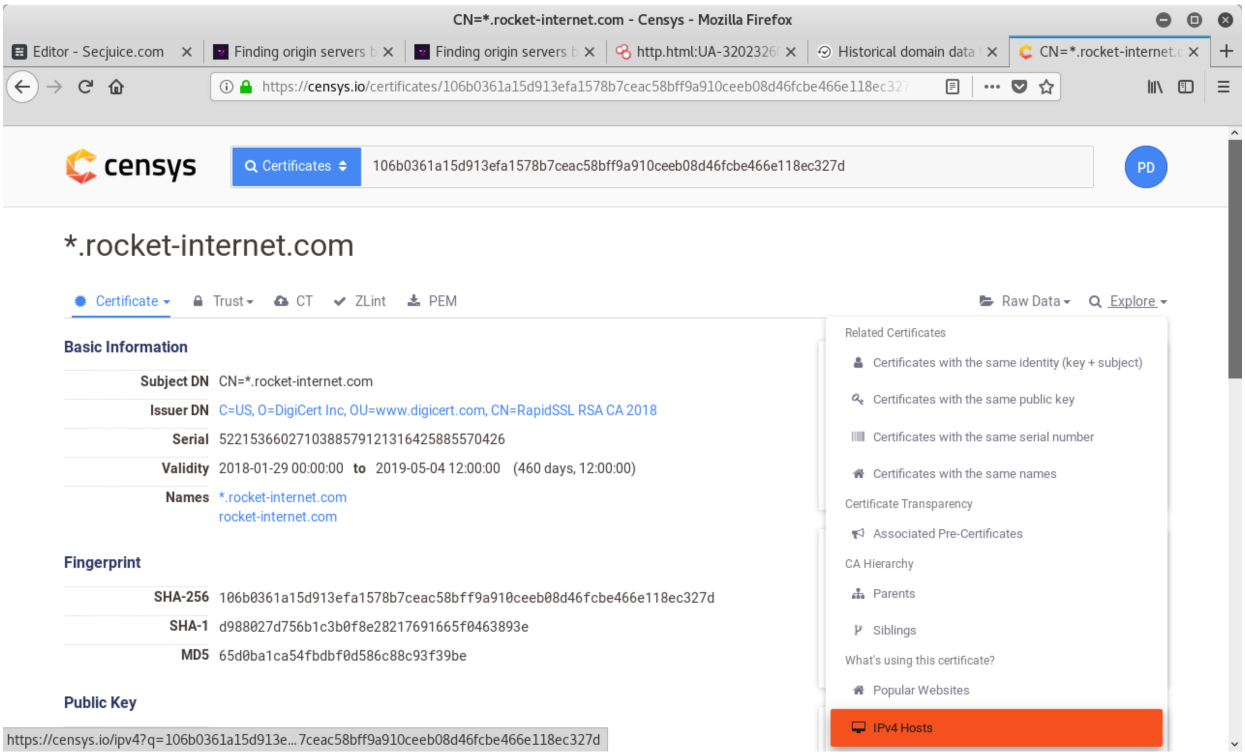

https://censys.io/certificates/106b0361a15d913efa1578b7ceac58bff9a910ceeb08d46fcbe466e118ec327

106b0361a15d913efa1578b7ceac58bff9a910ceeb08d46fcbe466e118ec327d

*.rocket-internet.com

Basic Information

Subject DN	CN=*.rocket-internet.com
Issuer DN	C=US, O=DigiCert Inc, OU=www.digicert.com, CN=RapidSSL RSA CA 2018
Serial	5221536602710388579121316425885570426
Validity	2018-01-29 00:00:00 to 2019-05-04 12:00:00 (460 days, 12:00:00)
Names	*.rocket-internet.com rocket-internet.com

Fingerprint

SHA-256	106b0361a15d913efa1578b7ceac58bff9a910ceeb08d46fcbe466e118ec327d
SHA-1	d988027d756b1c3b0f8e28217691665f0463893e
MD5	65d0ba1ca54fbdbf0d586c88c93f39be

Public Key

https://censys.io/ipv4?q=106b0361a15d913e...7ceac58bff9a910ceeb08d46fcbe466e118ec327d

You'll be presented a list of IPv4 Hosts using the specific certificate. One of those could be the origin IP.

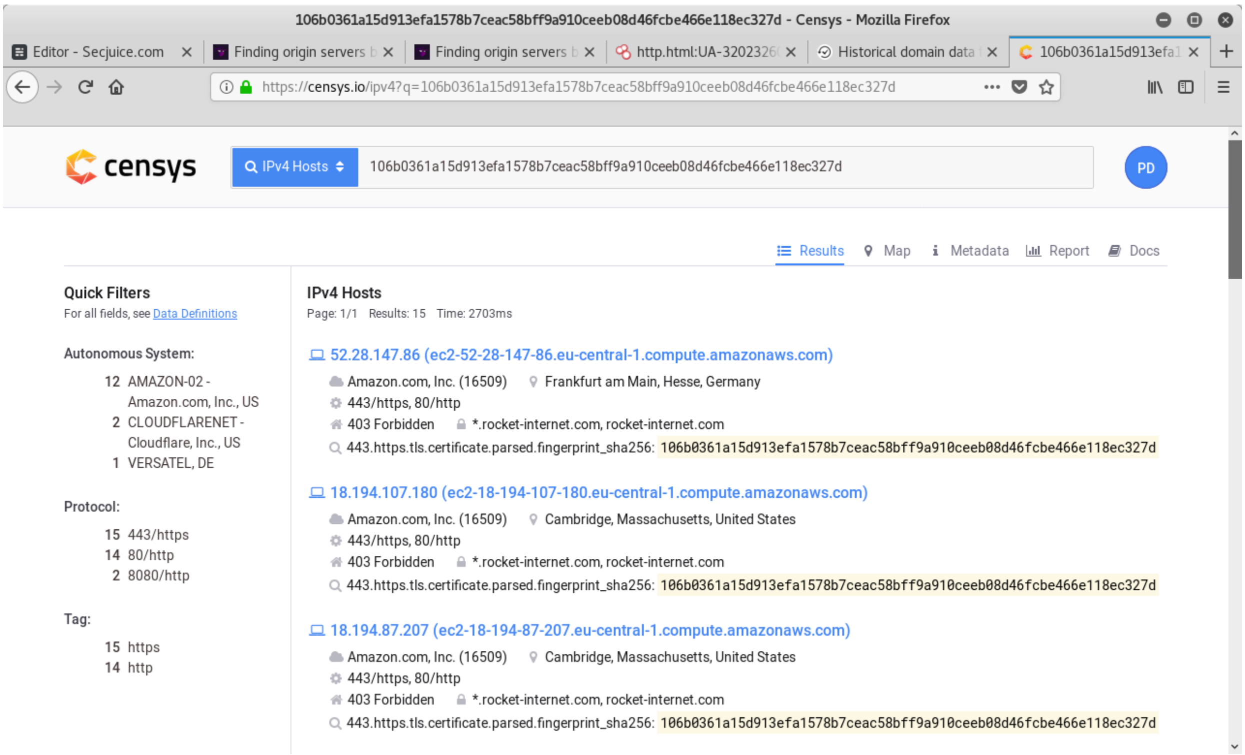
106b0361a15d913efa1578b7ceac58bff9a910ceeb08d46fcbe466e118ec327d - Censys - Mozilla Firefox
Editor - Secjuice.com
Finding origin servers b
Finding origin servers b
http.html:UA-3202326
Historical domain data
106b0361a15d913efa1
https://censys.io/ipv4?q=106b0361a15d913efa1578b7ceac58bff9a910ceeb08d46fcbe466e118ec327d
censys
IPv4 Hosts
106b0361a15d913efa1578b7ceac58bff9a910ceeb08d46fcbe466e118ec327d
PD
Results
Map
Metadata
Report
Docs

Quick Filters

For all fields, see Data Definitions

Autonomous System:

12 AMAZON-02 - Amazon.com, Inc., US
2 CLOUDFLARENET - Cloudflare, Inc., US
1 VERSATEL, DE

Protocol:

15 443/https
14 80/http
2 8080/http

Tag:

15 https
14 http

IPv4 Hosts

Page: 1/1 Results: 15 Time: 2703ms

52.28.147.86 (ec2-52-28-147-86.eu-central-1.compute.amazonaws.com)
Amazon.com, Inc. (16509) Frankfurt am Main, Hesse, Germany
443/https, 80/http
403 Forbidden *.rocket-internet.com, rocket-internet.com
443.https.tls.certificate.parsed.fingerprint_sha256: 106b0361a15d913efa1578b7ceac58bff9a910ceeb08d46fcbe466e118ec327d

18.194.107.180 (ec2-18-194-107-180.eu-central-1.compute.amazonaws.com)
Amazon.com, Inc. (16509) Cambridge, Massachusetts, United States
443/https, 80/http
403 Forbidden *.rocket-internet.com, rocket-internet.com
443.https.tls.certificate.parsed.fingerprint_sha256: 106b0361a15d913efa1578b7ceac58bff9a910ceeb08d46fcbe466e118ec327d

18.194.87.207 (ec2-18-194-87-207.eu-central-1.compute.amazonaws.com)
Amazon.com, Inc. (16509) Cambridge, Massachusetts, United States
443/https, 80/http
403 Forbidden *.rocket-internet.com, rocket-internet.com
443.https.tls.certificate.parsed.fingerprint_sha256: 106b0361a15d913efa1578b7ceac58bff9a910ceeb08d46fcbe466e118ec327d

You can verify by navigating to the IPs on port 443. Does it redirect to xyz123boot.com? Does it show the website directly on the IP?

1.2 Using a given SSL certificate

You are the FBI and want to shut down a child porn hidden service available under cheesecp5vaogohv.onion. To do that, you need the origin IP so you can contact the host and possibly also hunt down the operator then by following the money for example.

The hidden service has an SSL certificate. Finding IPv4 Hosts that use the same certificate can be done by just pasting its SHA1 fingerprint (without the colons) into the Censys IPv4 Hosts search. A badly configured web server can easily be found with this method.

2. DNS records

After suffering from multiple attacks, you decided to start using CloudFlare. Data-driven services like Censys for DNS records **still have your old A records** pointing to your web servers IP address.

A platform doing exactly this is SecurityTrails. Just enter the website domain into the search field and press enter. The "Historical Data" can be found in the sidebar on the left side.

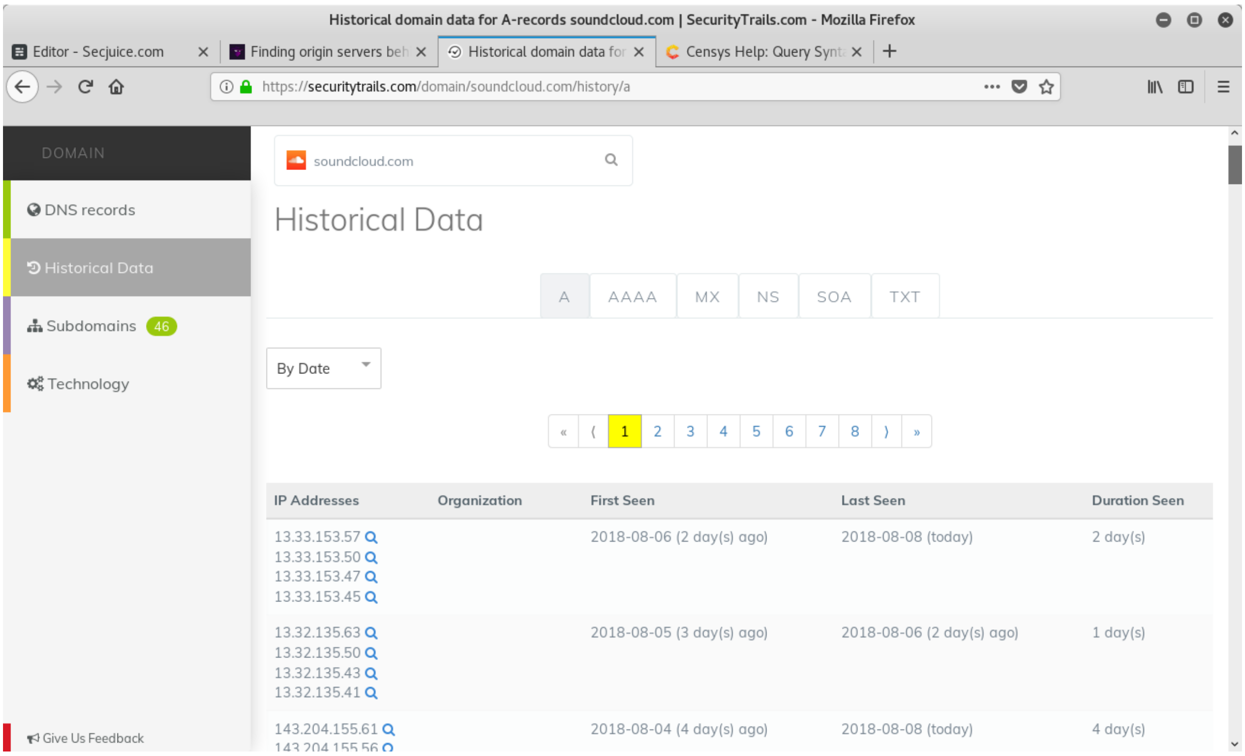

Historical domain data for A-records soundcloud.com | SecurityTrails.com - Mozilla Firefox
Editor - Secjuice.com
Finding origin servers beh
Historical domain data for
Censys Help: Query Synt
https://securitytrails.com/domain/soundcloud.com/history/a
DOMAIN
DNS records
Historical Data
Subdomains 46
Technology
Give Us Feedback
soundcloud.com
Historical Data
A
AAAA
MX
NS
SOA
TXT
By Date
« (1 2 3 4 5 6 7 8) »

IP Addresses	Organization	First Seen	Last Seen	Duration Seen
13.33.153.57 13.33.153.50 13.33.153.47 13.33.153.45		2018-08-06 (2 day(s) ago)	2018-08-08 (today)	2 day(s)
13.32.135.63 13.32.135.50 13.32.135.43 13.32.135.41		2018-08-05 (3 day(s) ago)	2018-08-06 (2 day(s) ago)	1 day(s)
143.204.155.61 143.204.155.56		2018-08-04 (4 day(s) ago)	2018-08-08 (today)	4 day(s)

Besides the old A records, even **current DNS records can leak the origin servers IP**. MX records, for example, are a common way of finding your IP. If the website is hosting its own mail server on the same server and IP as the web server, the origin server IP will be in the MX records.

3. HTTP Headers

3. HTTP Headers

With data-driven platforms that let anyone do powerful searches across a huge amount of data, even finding origin servers by comparing HTTP headers is a possibility.

Especially when having a pretty unique server header with various software including subversions, finding you is getting much easier.
This is also not limited to a single parameter. As mentioned in **1.1**, you can combine search parameters on Censys. The **likelihood of being found** with this method is **increasing with every less common header key or value** you are sending.

Let’s say you are sharing your server HTTP header with 1500 other web servers, which are sending the same header key and value combination. You are also using a new PHP framework sending a unique HTTP header (for example: X-Generated-Via: XYZ Framework). About 400 webmasters are using that framework in production yet. The intersection consists of a final amount of three servers. Going through those manually takes a few seconds and you found IP.

As an example, the search parameter at Censys for matching server headers is `80.http.get.headers.server:`.

Finding websites being served by CloudFlare works like this:

```
80.http.get.headers.server: cloudflare
```

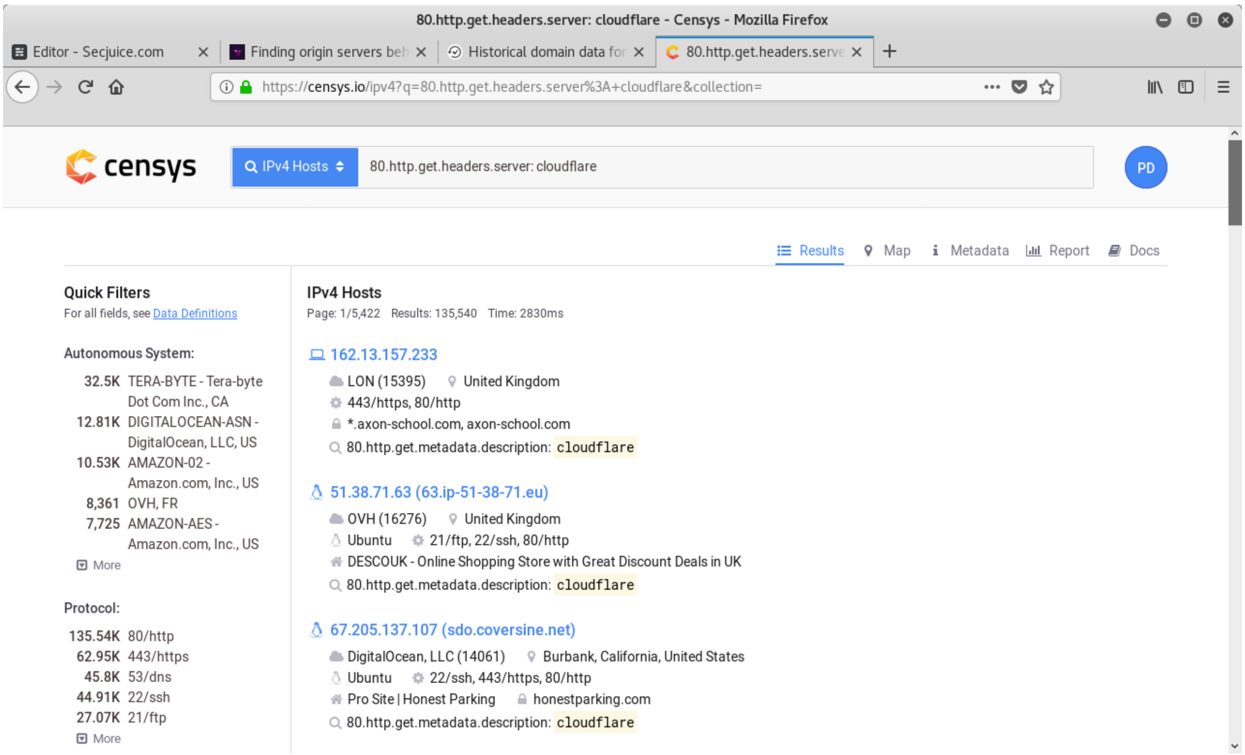
80.http.get.headers.server: cloudflare - Censys - Mozilla Firefox
Editor - Secjuice.com
Finding origin servers beh
Historical domain data for
80.http.get.headers.serve
censys
IPv4 Hosts
PD
Results
Map
Metadata
Report
Docs

https://censys.io/ipv4?q=80.http.get.headers.server%3A+cloudflare&collection=

80.http.get.headers.server: cloudflare

Quick Filters

For all fields, see Data Definitions

Autonomous System:

32.5K	TERA-BYTE - Tera-byte Dot Com Inc., CA
12.81K	DIGITALOCEAN-ASN - DigitalOcean, LLC, US
10.53K	AMAZON-02 - Amazon.com, Inc., US
8,361	OVH, FR
7,725	AMAZON-AES - Amazon.com, Inc., US

More

Protocol:

135.54K	80/http
62.95K	443/https
45.8K	53/dns
44.91K	22/ssh
27.07K	21/ftp

More

IPv4 Hosts

Page: 1/5,422 Results: 135,540 Time: 2830ms

162.13.157.233
LON (15395) United Kingdom
443/https, 80/http
*.axon-school.com, axon-school.com
80.http.get.metadata.description: cloudflare

51.38.71.63 (63.ip-51-38-71.eu)
OVH (16276) United Kingdom
Ubuntu 21/ftp, 22/ssh, 80/http
DESCOUK - Online Shopping Store with Great Discount Deals in UK
80.http.get.metadata.description: cloudflare

67.205.137.107 (sdo.coversine.net)
DigitalOcean, LLC (14061) Burbank, California, United States
Ubuntu 22/ssh, 443/https, 80/http
Pro Site | Honest Parking honestparking.com
80.http.get.metadata.description: cloudflare

4. Applications and Services

A TOR hidden service or a website being served through CloudFlare is a normal website. Starting a quick pentest could reveal the IP as well.

Headers like the HTTP server header can be used to find possible exploits for the services and versions in use. When gaining access to the server, you can obviously easily find the IP.

Another attempt would be to find edge cases triggering errors. Error messages can reveal sensitive information. Those pieces of information could be just the IP itself or anything that can be used as parameters for the other methods described here. This is all about being creative, doing recon and combining.

Running gobuster to find files and directories during the recon phase should be done in every pentest. Things you could find are logs, database dumps/backups and more.

Also, worth a check is to find out if you can make the application powering the website to interact with other services. If you're not the NSA, you probably can't get the IP if they are just consuming an API. But as an example, maybe you can set an avatar on the website and provide an URL to the picture instead of uploading it. If they are downloading it, they are probably doing it from their origin server. Now the IP is in your logs.

This was just a quick overview. You can do a large portion of all the pentest magic you would normally do and there are many mistakes webmasters could have made.

5. Content

In case the origin server IP is returning the content of the website as well, the massive amount of data searchable on the web got you again.

Going through the websites source code, you are looking for **unique pieces of code**. Third party services (e.g. Google Analytics, reCAPTCHA) with access/identifier keys in the JavaScript are a good start.

Example Google Analytics Tracking Code taken from HackTheBox website:

```
ga('create', 'UA-93577176-1', 'auto');
```

Filtering Censys data by the body/source can be done with the `80.http.get.body:` parameter.

Unfortunately, the normal search field has limitations. You can request research access at Censys, which allows you to do much more powerful queries via Google BigQuery.

Shodan, a service similar to Censys, provides a `http.html` search parameter, too.

An example search:

https://www.shodan.io/search?query=http.html%3AUA-32023260-1

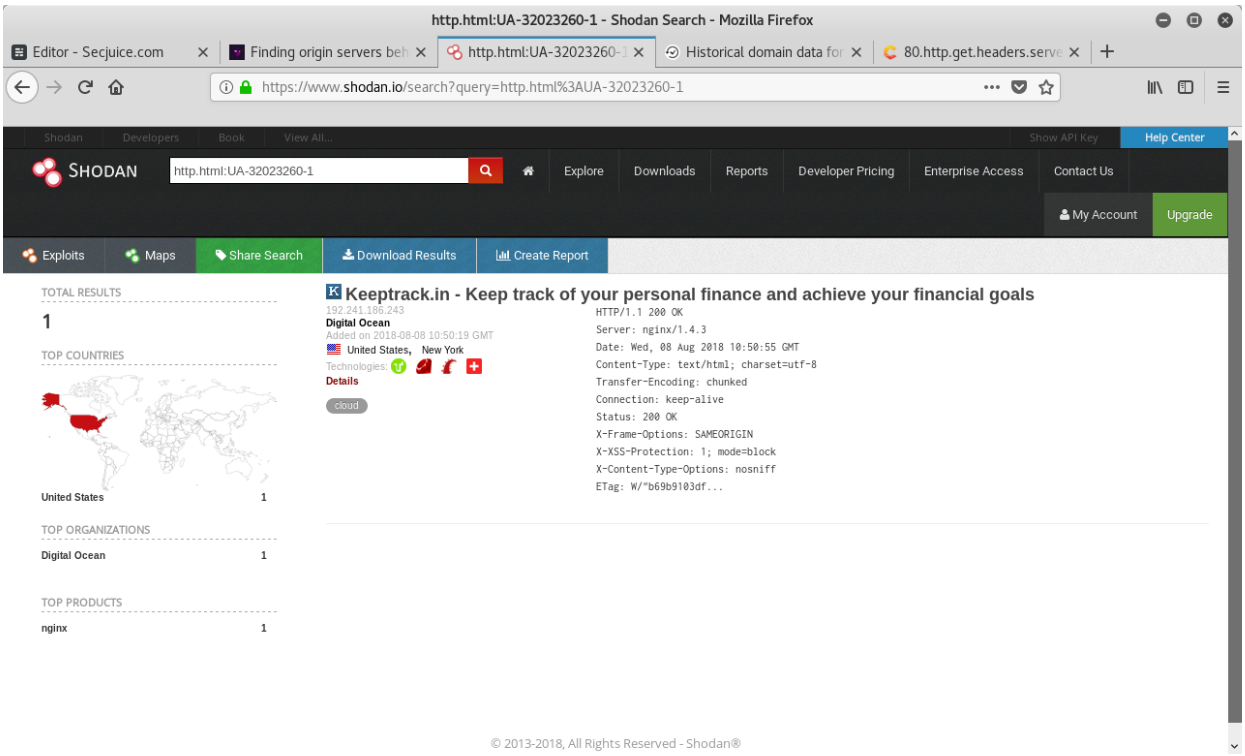

TOTAL RESULTS

1

TOP COUNTRIES

United States 1

TOP ORGANIZATIONS

Digital Ocean 1

TOP PRODUCTS

nginx 1

Keeptrack.in - Keep track of your personal finance and achieve your financial goals

192.241.186.243
Digital Ocean
Added on 2018-08-08 10:50:19 GMT
United States, New York
Technologies:
Details
cloud

```
HTTP/1.1 200 OK
Server: nginx/1.4.3
Date: Wed, 08 Aug 2018 10:50:55 GMT
Content-Type: text/html; charset=utf-8
Transfer-Encoding: chunked
Connection: keep-alive
Status: 200 OK
X-Frame-Options: SAMEORIGIN
X-XSS-Protection: 1; mode=block
X-Content-Type-Options: nosniff
ETag: W/"b69b9103df...
```

© 2013-2018, All Rights Reserved - Shodan®

Final notes

In the end, finding the origin IP behind TOR hidden services or reverse-proxy services like CloudFlare mostly **requires a certain amount of understanding of the web and creativity**.

Securing a Web Hidden Service by x0rz explains countermeasures for TOR hidden service operators against several methods covered in this article. It also contains glorious fails, in which hidden services didn't master opsec, so security researchers could unmask them.

Do you run a hidden service or are you using CloudFlare? Made any of the mistakes described above?

What other methods could reveal origin IPs? I'd be glad to hear about them so we can make this an even more comprehensive resource.

Main Image Credit : The awesome piece of artwork used to head this article is called 'Mystic Cat' and it was created by graphic designer Alexa Erkaeva.

Update Note (19/8/2018) in response to Higinio “w0rmer” Ochoa (CloudFlare Security Engineer)

Shortly after publishing the article, a Security Engineer at CloudFlare added a couple of valid comments. Kudos to CloudFlare as their security team seems to be up to date in terms of new articles and everything related to security especially regarding their own service.

> *Oooh looks like someone hasn’t heard of Argo Tunnels https://t.co/aVWJBMX4N5*
>
> — Higinio “w0rmer” Ochoa (@0x686967) 19. August 2018

All examples in this article work like this when making the mistakes de