Breaking the bad, Stopping the ugly By using Open Source

Ulrika Vincent
dnstapir.se

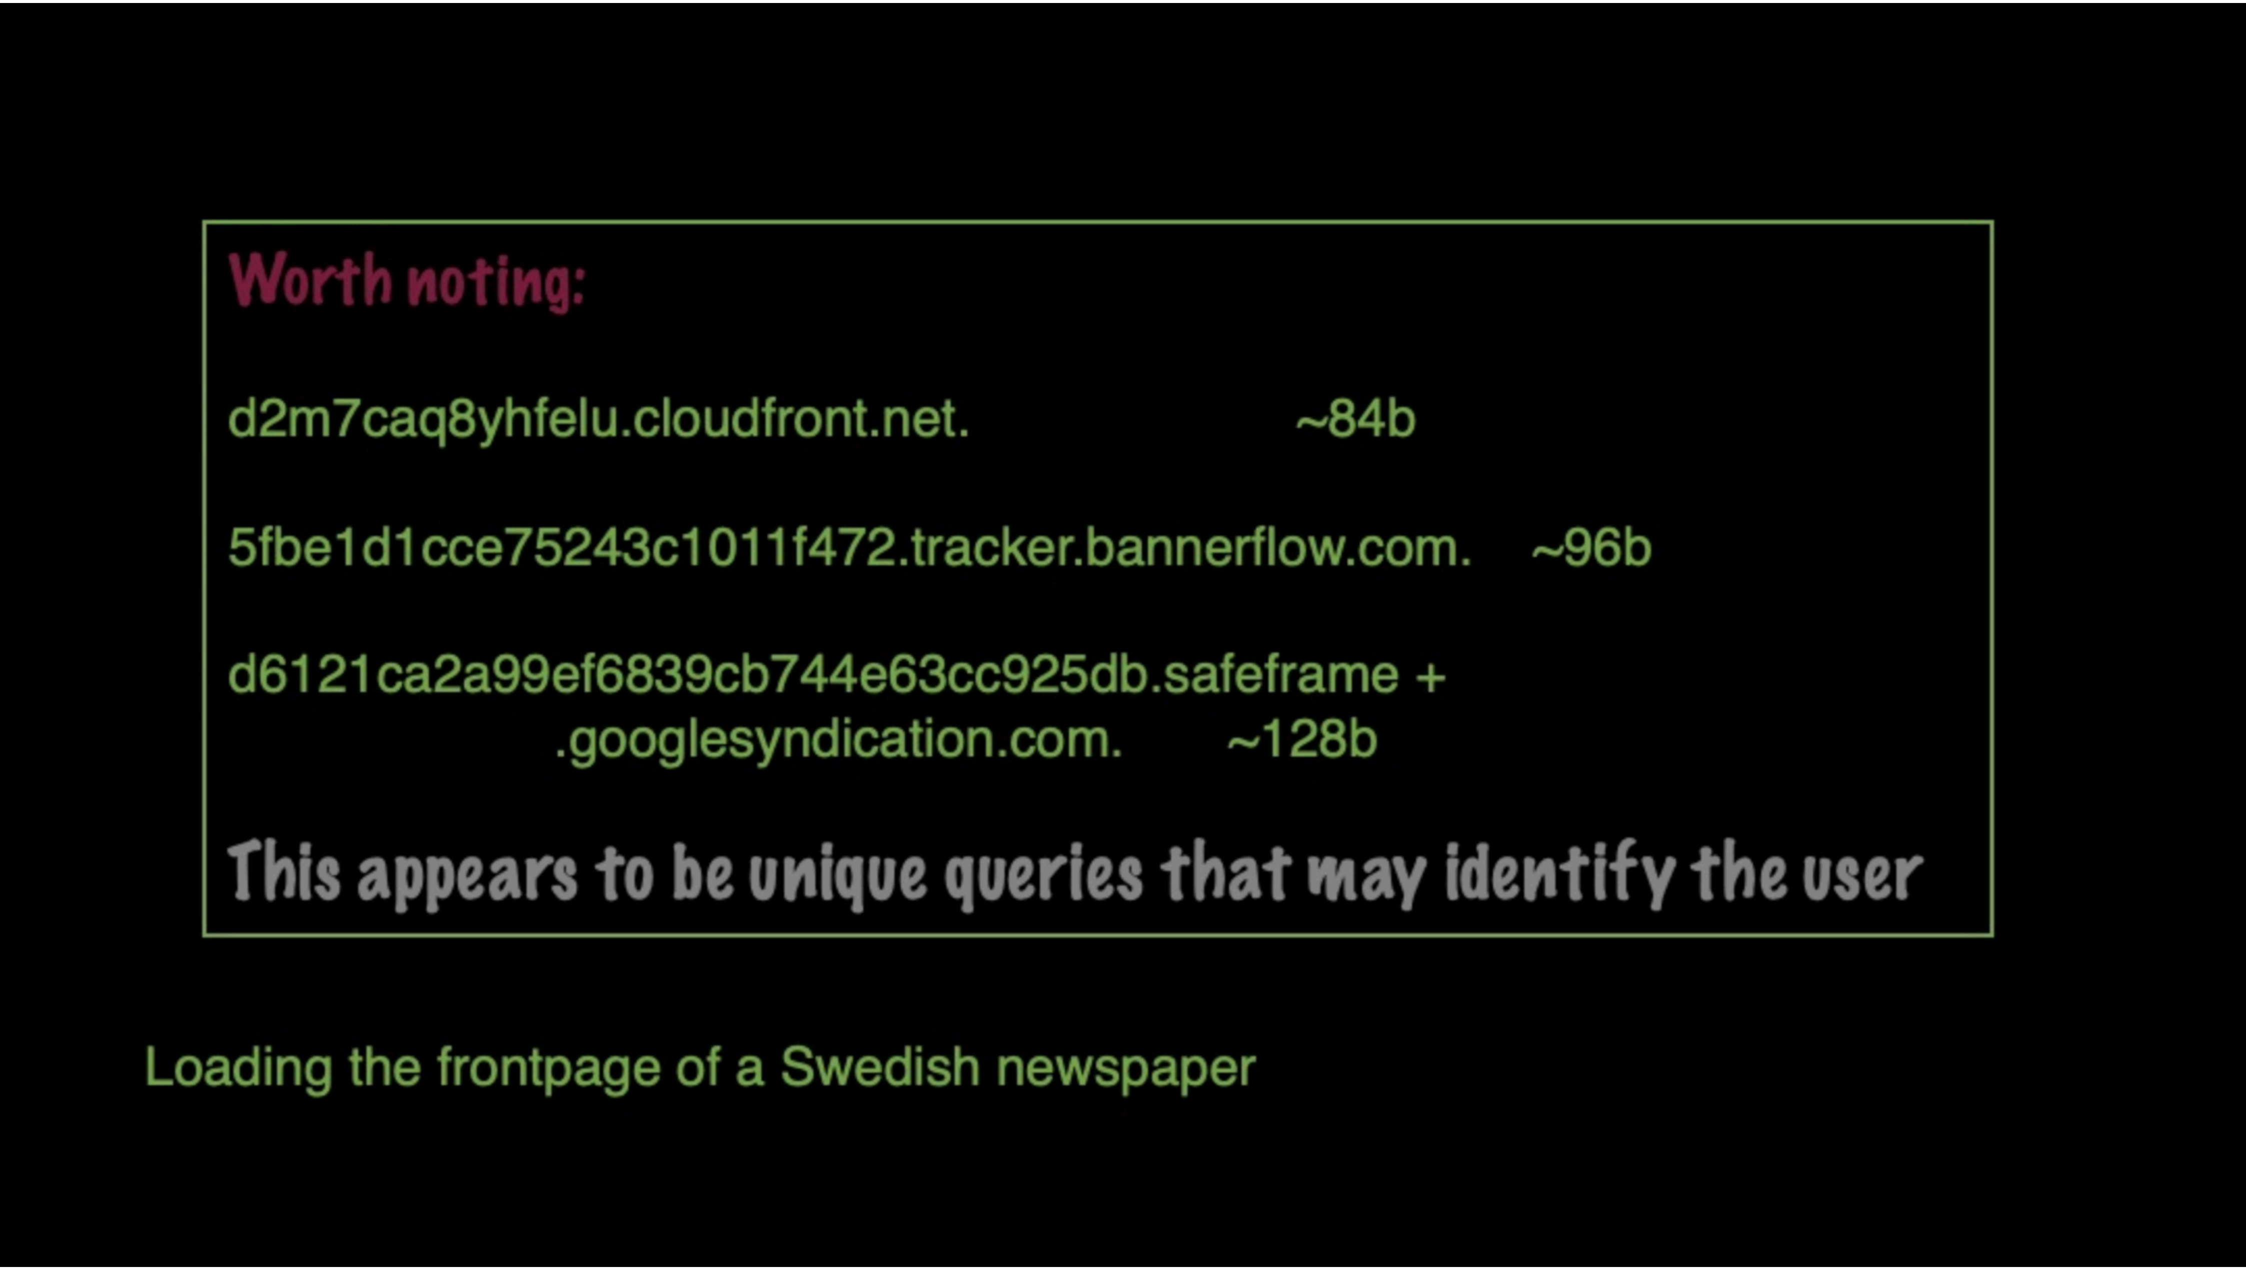

Worth noting:

d2m7caq8yhfelu.cloudfront.net.	~84b
5fbe1d1cce75243c1011f472.tracker.bannerflow.com.	~96b
d6121ca2a99ef6839cb744e63cc925db.safeframe + .googlesyndication.com.	~128b

This appears to be unique queries that may identify the user

Loading the frontpage of a Swedish newspaper

Some TAPIRs

DNS TAPIR

A privacy first,
open source,
local decisions
and open data
DNS query analytics platform

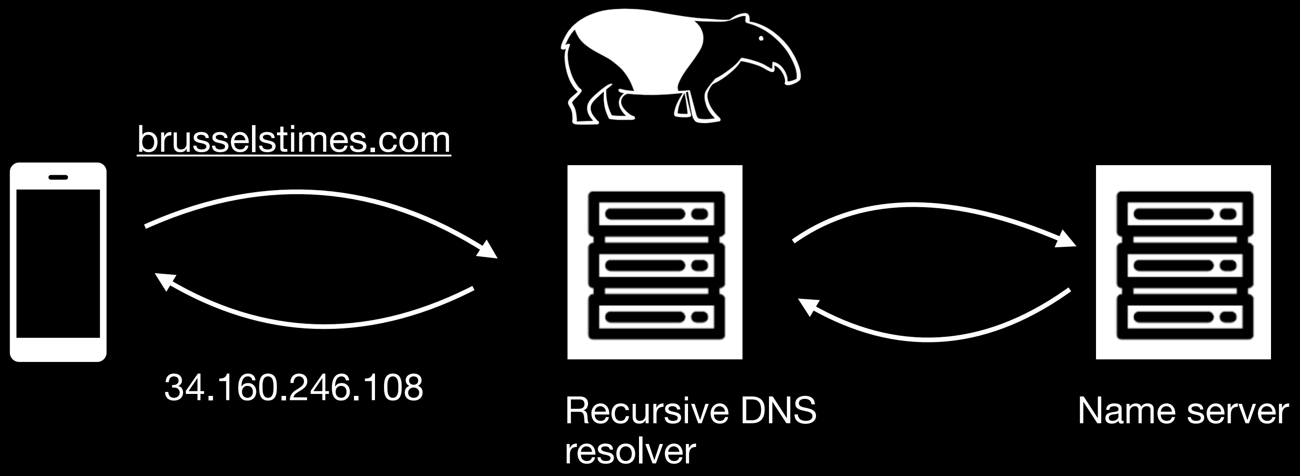

DNS TAPIR

A privacy first,
open source,
local decisions
and open data
DNS query analytics platform

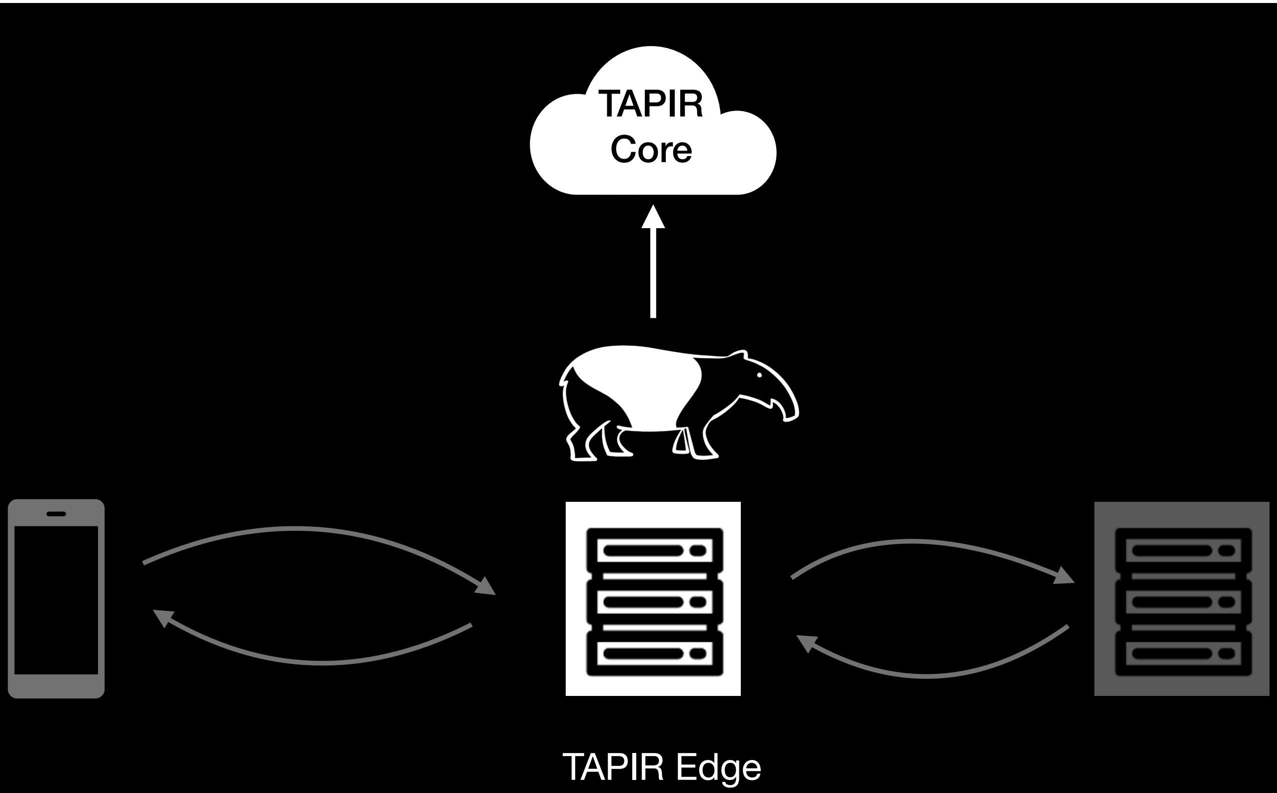

DNS TAPIR

A privacy first,
open source,
local decisions
and open data
DNS query analytics platform

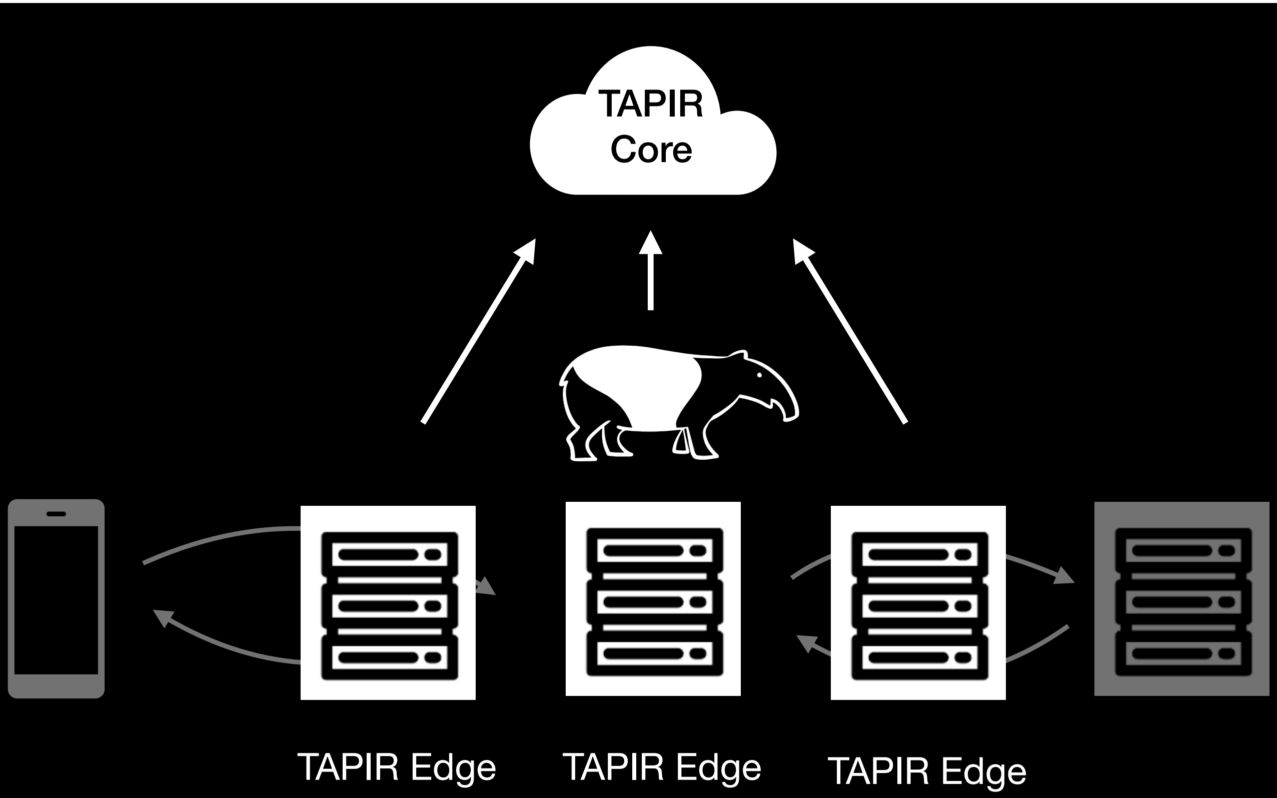

Building a data commons

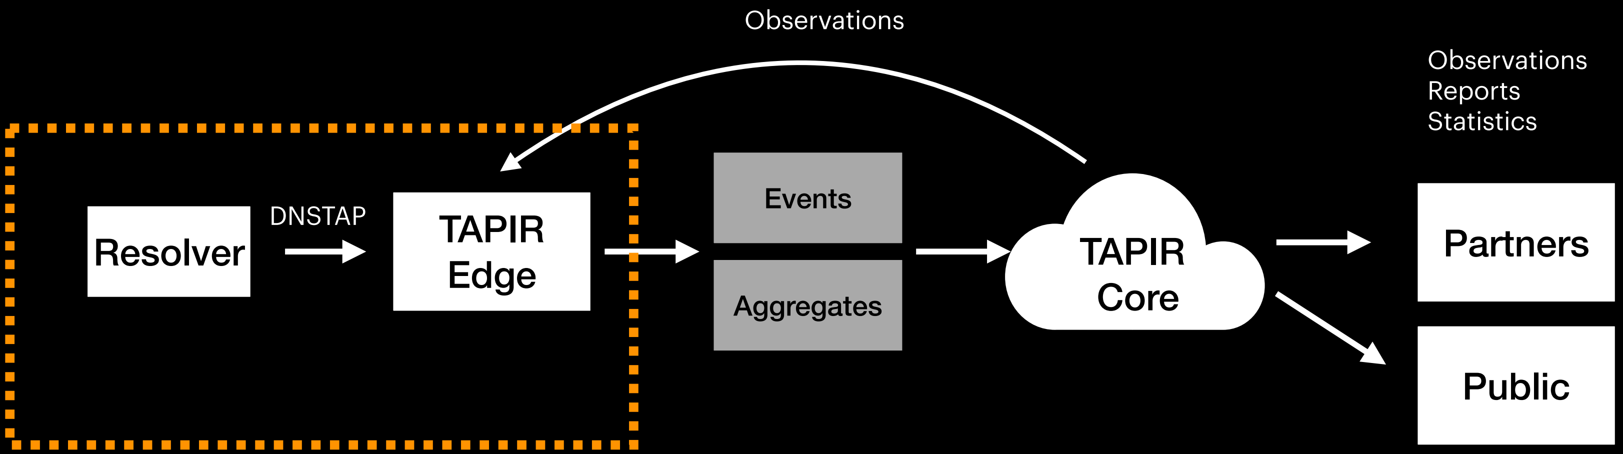

The main purpose of DNS TAPIR is to make DNS data transparent and available to interested parties by addressing the challenge of it being highly privacy sensitive.

We aim to create a data commons.

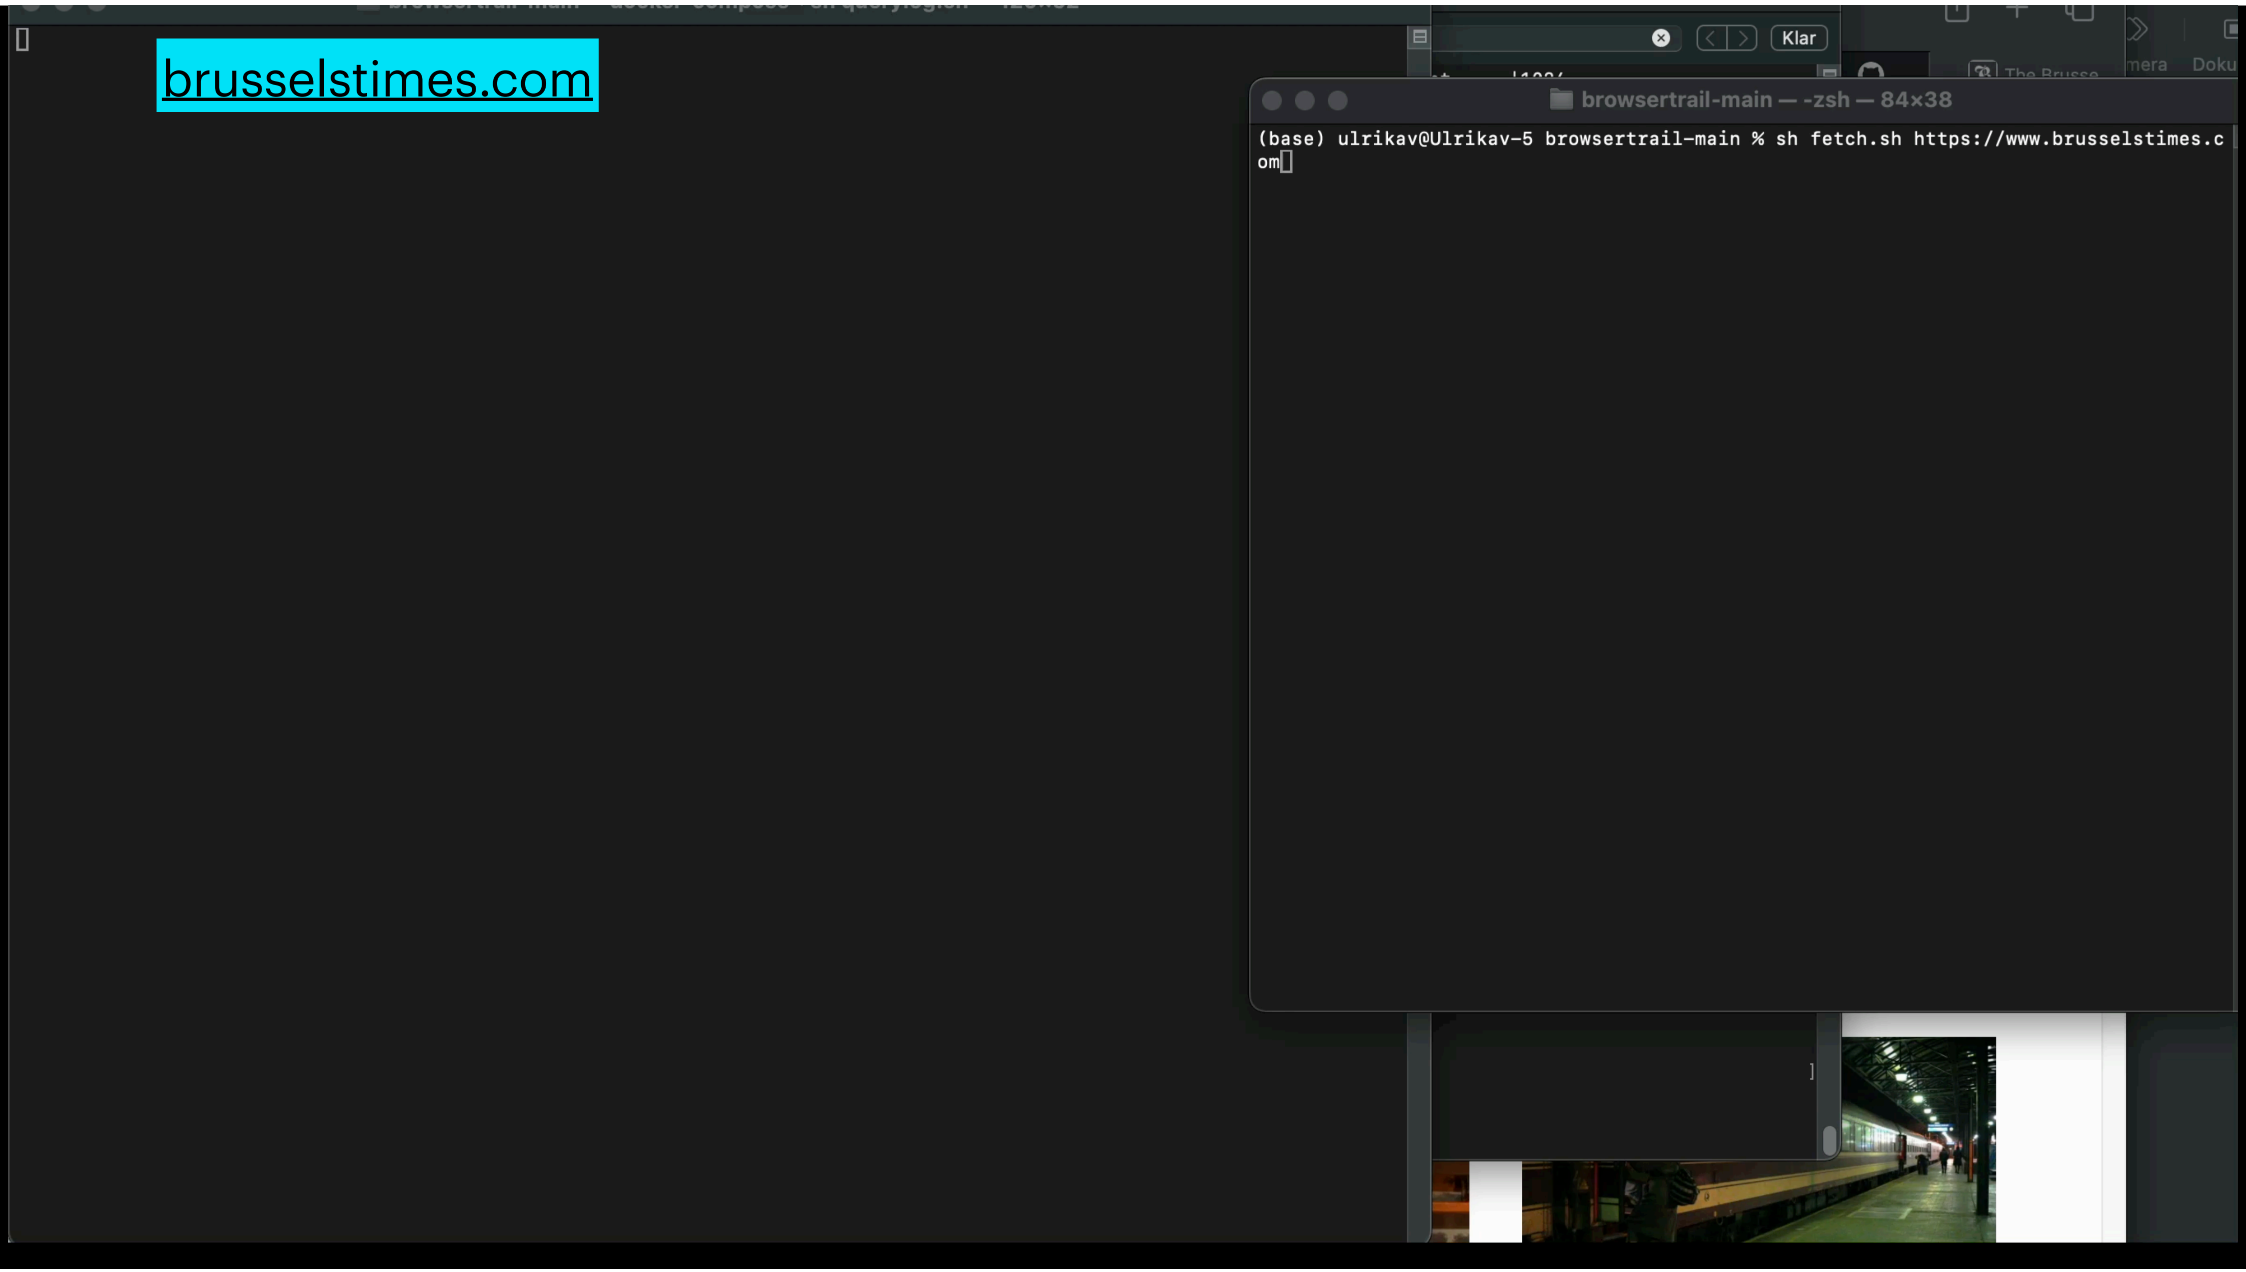

brusselstimes.com

```
(base) ulrikav@Ulrikav-5 browsertrail-main % sh fetch.sh https://www.brusselstimes.c
om
```

brusselstimes.com

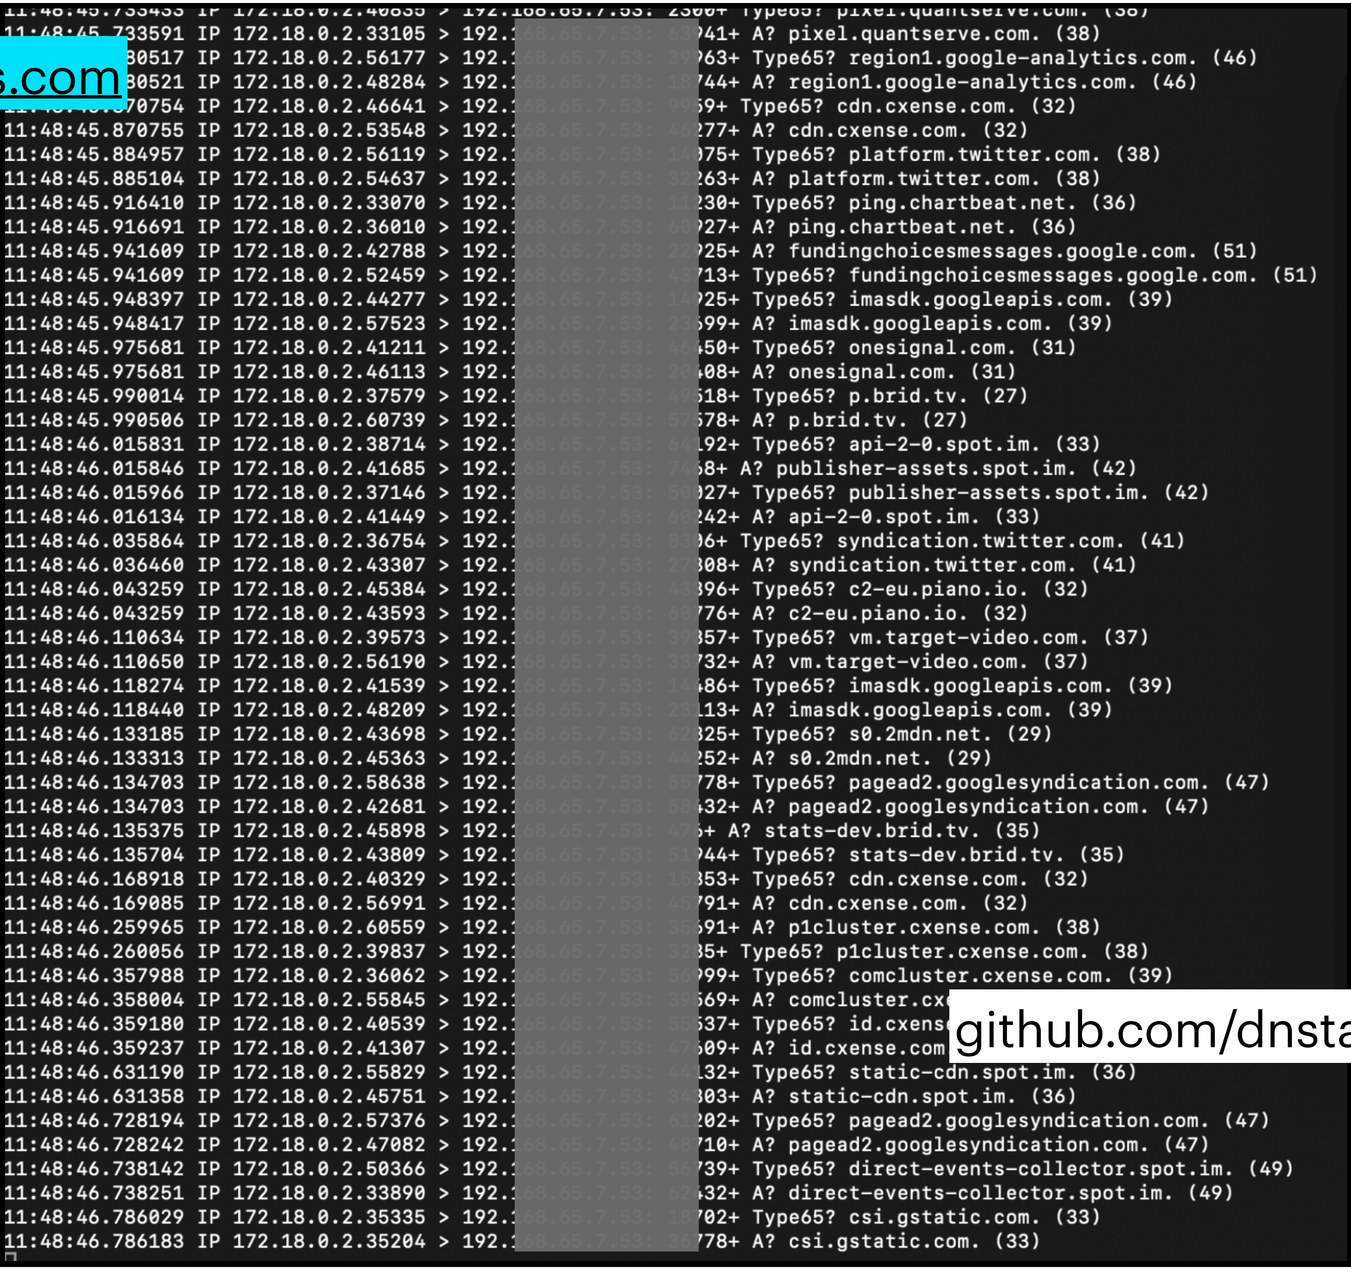

```
11:48:45.733433 IP 172.18.0.2.40835 > 192.168.65.7.53: 2300+ Type65? pixel.quantserve.com. (38)
11:48:45.733591 IP 172.18.0.2.33105 > 192.1[redacted]41+ A? pixel.quantserve.com. (38)
          30517 IP 172.18.0.2.56177 > 192.1[redacted]63+ Type65? region1.google-analytics.com. (46)
          30521 IP 172.18.0.2.48284 > 192.1[redacted]44+ A? region1.google-analytics.com. (46)
           70754 IP 172.18.0.2.46641 > 192.1[redacted]59+ Type65? cdn.cxense.com. (32)
11:48:45.870755 IP 172.18.0.2.53548 > 192.1[redacted]277+ A? cdn.cxense.com. (32)
11:48:45.884957 IP 172.18.0.2.56119 > 192.1[redacted]75+ Type65? platform.twitter.com. (38)
11:48:45.885104 IP 172.18.0.2.54637 > 192.1[redacted]263+ A? platform.twitter.com. (38)
11:48:45.916410 IP 172.18.0.2.33070 > 192.1[redacted]230+ Type65? ping.chartbeat.net. (36)
11:48:45.916691 IP 172.18.0.2.36010 > 192.1[redacted]27+ A? ping.chartbeat.net. (36)
11:48:45.941609 IP 172.18.0.2.42788 > 192.1[redacted]25+ A? fundingchoicesmessages.google.com. (51)
11:48:45.941609 IP 172.18.0.2.52459 > 192.1[redacted]13+ Type65? fundingchoicesmessages.google.com. (51)
11:48:45.948397 IP 172.18.0.2.44277 > 192.1[redacted]25+ Type65? imasdk.googleapis.com. (39)
11:48:45.948417 IP 172.18.0.2.57523 > 192.1[redacted]99+ A? imasdk.googleapis.com. (39)
11:48:45.975681 IP 172.18.0.2.41211 > 192.1[redacted]50+ Type65? onesignal.com. (31)
11:48:45.975681 IP 172.18.0.2.46113 > 192.1[redacted]08+ A? onesignal.com. (31)
11:48:45.990014 IP 172.18.0.2.37579 > 192.1[redacted]18+ Type65? p.brid.tv. (27)
11:48:45.990506 IP 172.18.0.2.60739 > 192.1[redacted]78+ A? p.brid.tv. (27)
11:48:46.015831 IP 172.18.0.2.38714 > 192.1[redacted]192+ Type65? api-2-0.spot.im. (33)
11:48:46.015846 IP 172.18.0.2.41685 > 192.1[redacted]8+ A? publisher-assets.spot.im. (42)
11:48:46.015966 IP 172.18.0.2.37146 > 192.1[redacted]27+ Type65? publisher-assets.spot.im. (42)
11:48:46.016134 IP 172.18.0.2.41449 > 192.1[redacted]242+ A? api-2-0.spot.im. (33)
11:48:46.035864 IP 172.18.0.2.36754 > 192.1[redacted]6+ Type65? syndication.twitter.com. (41)
11:48:46.036460 IP 172.18.0.2.43307 > 192.1[redacted]308+ A? syndication.twitter.com. (41)
11:48:46.043259 IP 172.18.0.2.45384 > 192.1[redacted]396+ Type65? c2-eu.piano.io. (32)
11:48:46.043259 IP 172.18.0.2.43593 > 192.1[redacted]776+ A? c2-eu.piano.io. (32)
11:48:46.110634 IP 172.18.0.2.39573 > 192.1[redacted]357+ Type65? vm.target-video.com. (37)
11:48:46.110650 IP 172.18.0.2.56190 > 192.1[redacted]732+ A? vm.target-video.com. (37)
11:48:46.118274 IP 172.18.0.2.41539 > 192.1[redacted]86+ Type65? imasdk.googleapis.com. (39)
11:48:46.118440 IP 172.18.0.2.48209 > 192.1[redacted]113+ A? imasdk.googleapis.com. (39)
11:48:46.133185 IP 172.18.0.2.43698 > 192.1[redacted]325+ Type65? s0.2mdn.net. (29)
11:48:46.133313 IP 172.18.0.2.45363 > 192.1[redacted]252+ A? s0.2mdn.net. (29)
11:48:46.134703 IP 172.18.0.2.58638 > 192.1[redacted]778+ Type65? pagead2.googlesyndication.com. (47)
11:48:46.134703 IP 172.18.0.2.42681 > 192.1[redacted]32+ A? pagead2.googlesyndication.com. (47)
11:48:46.135375 IP 172.18.0.2.45898 > 192.1[redacted]+ A? stats-dev.brid.tv. (35)
11:48:46.135704 IP 172.18.0.2.43809 > 192.1[redacted]44+ Type65? stats-dev.brid.tv. (35)
11:48:46.168918 IP 172.18.0.2.40329 > 192.1[redacted]353+ Type65? cdn.cxense.com. (32)
11:48:46.169085 IP 172.18.0.2.56991 > 192.1[redacted]791+ A? cdn.cxense.com. (32)
11:48:46.259965 IP 172.18.0.2.60559 > 192.1[redacted]91+ A? p1cluster.cxense.com. (38)
11:48:46.260056 IP 172.18.0.2.39837 > 192.1[redacted]35+ Type65? p1cluster.cxense.com. (38)
11:48:46.357988 IP 172.18.0.2.36062 > 192.1[redacted]99+ Type65? comcluster.cxense.com. (39)
11:48:46.358004 IP 172.18.0.2.55845 > 192.1[redacted]69+ A? comcluster.cx
11:48:46.359180 IP 172.18.0.2.40539 > 192.1[redacted]37+ Type65? id.cxense
11:48:46.359237 IP 172.18.0.2.41307 > 192.1[redacted]09+ A? id.cxense.com
11:48:46.631190 IP 172.18.0.2.55829 > 192.1[redacted]132+ Type65? static-cdn.spot.im. (36)
11:48:46.631358 IP 172.18.0.2.45751 > 192.1[redacted]303+ A? static-cdn.spot.im. (36)
11:48:46.728194 IP 172.18.0.2.57376 > 192.1[redacted]202+ Type65? pagead2.googlesyndication.com. (47)
11:48:46.728242 IP 172.18.0.2.47082 > 192.1[redacted]710+ A? pagead2.googlesyndication.com. (47)
11:48:46.738142 IP 172.18.0.2.50366 > 192.1[redacted]739+ Type65? direct-events-collector.spot.im. (49)
11:48:46.738251 IP 172.18.0.2.33890 > 192.1[redacted]432+ A? direct-events-collector.spot.im. (49)
11:48:46.786029 IP 172.18.0.2.35335 > 192.1[redacted]702+ Type65? csi.gstatic.com. (33)
11:48:46.786183 IP 172.18.0.2.35204 > 192.1[redacted]778+ A? csi.gstatic.com. (33)
```

github.com/dnstapir/BrowserTrail

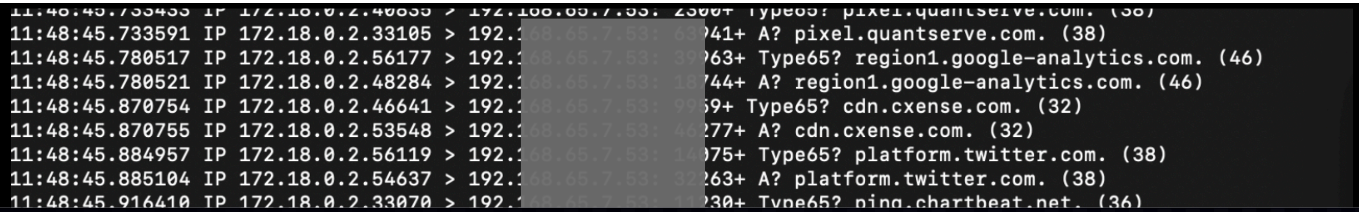

```
11:48:45.733433 IP 172.18.0.2.40835 > 192.168.65.7.53: 2360+ Type65? pixel.quantserve.com. (38)
11:48:45.733591 IP 172.18.0.2.33105 > 192.1[redacted]41+ A? pixel.quantserve.com. (38)
11:48:45.780517 IP 172.18.0.2.56177 > 192.1[redacted]63+ Type65? region1.google-analytics.com. (46)
11:48:45.780521 IP 172.18.0.2.48284 > 192.1[redacted]44+ A? region1.google-analytics.com. (46)
11:48:45.870754 IP 172.18.0.2.46641 > 192.1[redacted]59+ Type65? cdn.cxense.com. (32)
11:48:45.870755 IP 172.18.0.2.53548 > 192.1[redacted]277+ A? cdn.cxense.com. (32)
11:48:45.884957 IP 172.18.0.2.56119 > 192.1[redacted]75+ Type65? platform.twitter.com. (38)
11:48:45.885104 IP 172.18.0.2.54637 > 192.1[redacted]263+ A? platform.twitter.com. (38)
11:48:45.916410 IP 172.18.0.2.33070 > 192.1[redacted]230+ Type65? ping.chartbeat.net. (36)
```

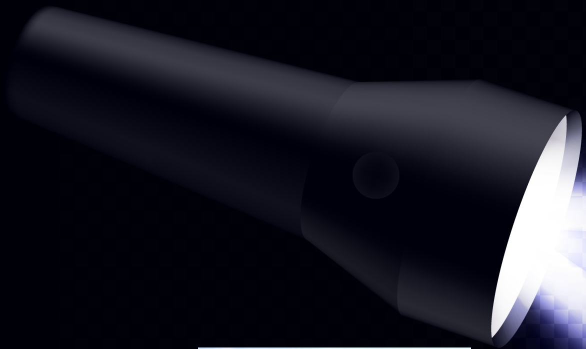

Examples of interesting (old) domains:

google—analytics.com
google-anallytics.com

org.naijaisblessed.sitkey.id.online.upgrade.system.update.new.com.bankofamerica.

net.eu-app.user-eu5.myapple-unlock.cgi-key.auth.id.confirm.en-gb.com.apple.support.

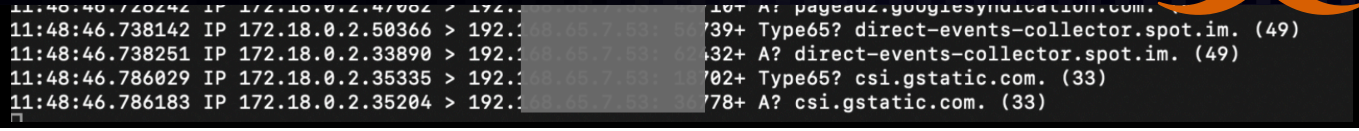

```
11:48:46.728242 IP 172.18.0.2.47082 > 192.1[redacted]16+ A? pagead2.googlesyndication.com.
11:48:46.738142 IP 172.18.0.2.50366 > 192.1[redacted]739+ Type65? direct-events-collector.spot.im. (49)
11:48:46.738251 IP 172.18.0.2.33890 > 192.1[redacted]32+ A? direct-events-collector.spot.im. (49)
11:48:46.786029 IP 172.18.0.2.35335 > 192.1[redacted]702+ Type65? csi.gstatic.com. (33)
11:48:46.786183 IP 172.18.0.2.35204 > 192.1[redacted]778+ A? csi.gstatic.com. (33)
```

Spying on DNS users is a BAD idea!

Image: https://webhostinggeeks.com/guides/privacy/

What makes TAPIR different?

Common solutions

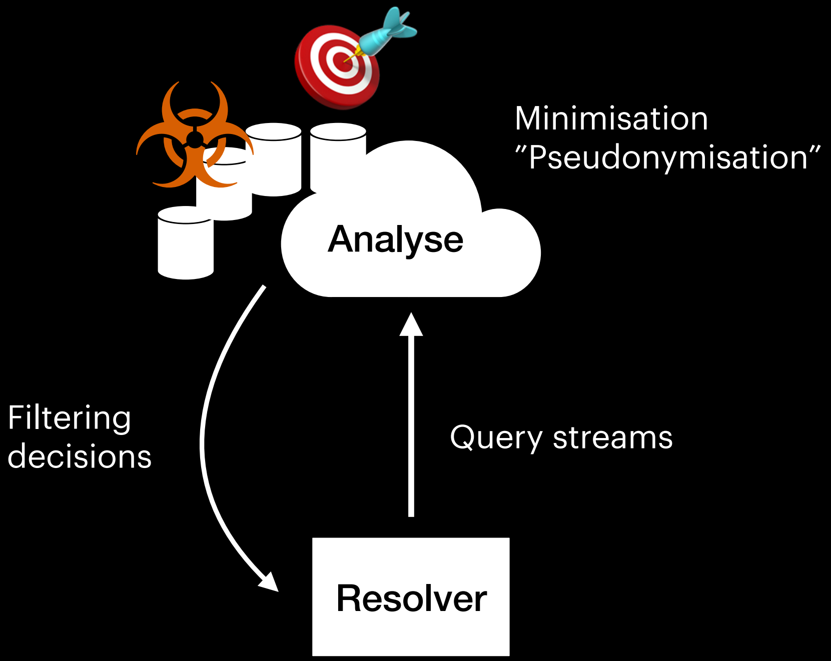
Minimisation
"Pseudonymisation"

Analyse

Filtering decisions

Query streams

Resolver

Centralised storage of browsing behaviour data

Lack of transparency into filtering decisions

No or very little open data

DNS TAPIR

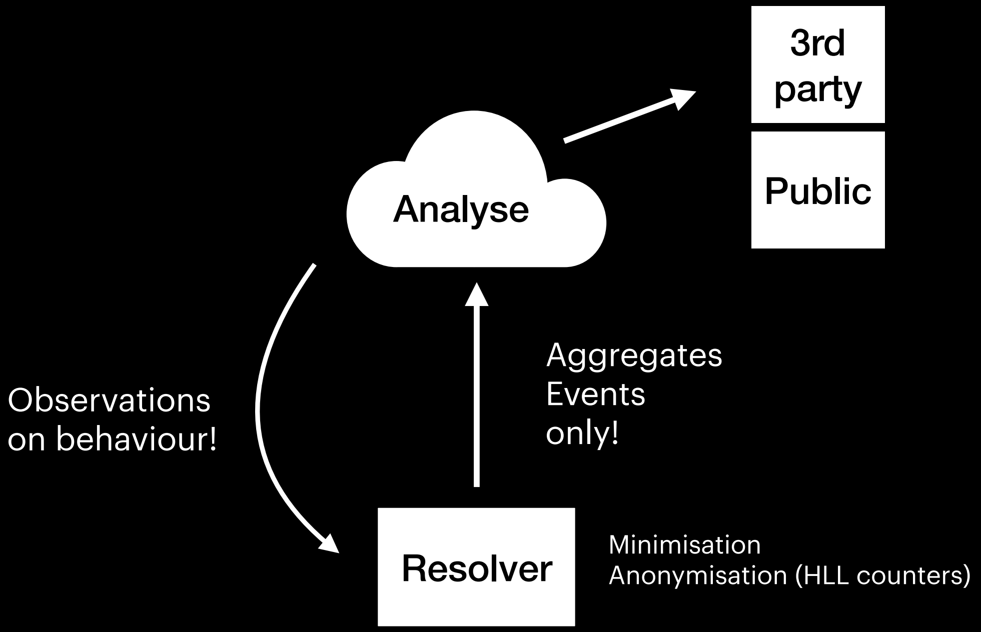
3rd party

Public

Analyse

Observations on behaviour!

Aggregates
Events
only!

Resolver

Minimisation
Anonymisation (HLL counters)

NO sensitive query data leaves resolver

Full local control over filtering

Aims to create a data commons

Common products

🤑

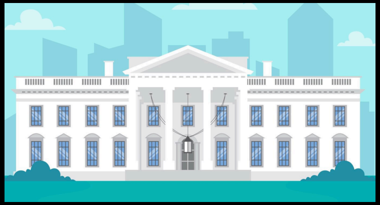

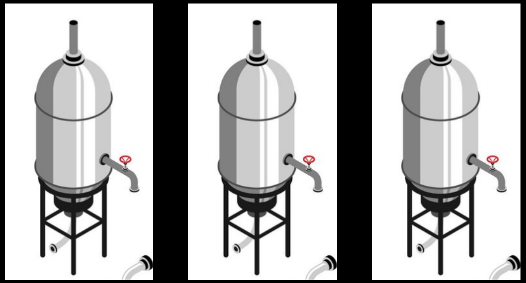

Driven by profits - influenced by politics - hidden in silos

DNS TAPIR

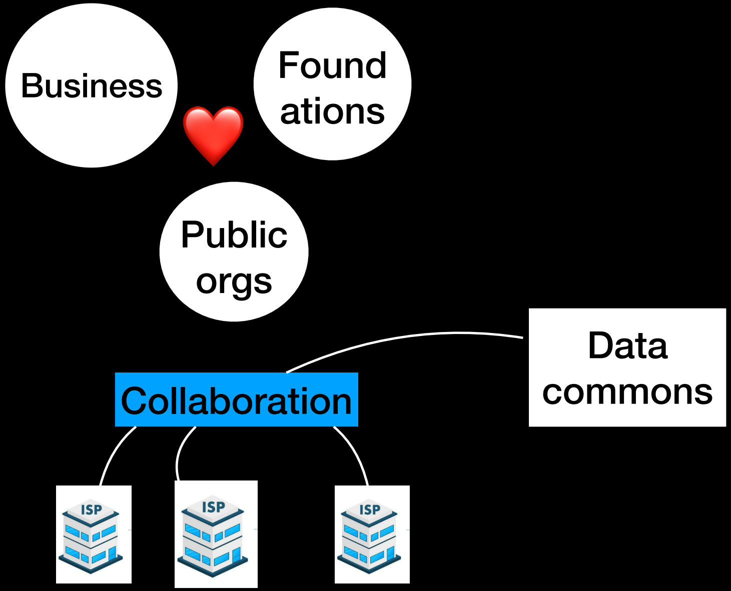

Business

Found ations

Public orgs

Collaboration

Data commons

Funding - Balance of interests - shared results

https://www.freepik.com/free-vector/white-house-illustration-flat-design_12552702.htm#fromView=search&page=1&position=10&uuid=b88dc291-70a9-4f25-afc2-82cfcee558e3&query=government

https://www.vecteezy.com/vector-art/29899491-isp-internet-service-provider-company-that-provides-web-access-vector-stock-illustration

https://www.cfostrategiesllc.com/industries/non-profits/

Problems with common solutions:

- pseudonymisation doesn’t work

- aggregation, unless done right, doesn’t work

- centralized storage of sensitive data becomes a target

- ”only authorized analysts” is an illusion.

Better solution:

- Open data commons - leads to better privacy by design

- Minimisation at source

- Distributed, local storage, more difficult to attack

- Differential privacy (aggregation done right)

Data you don’t have can’t be lost

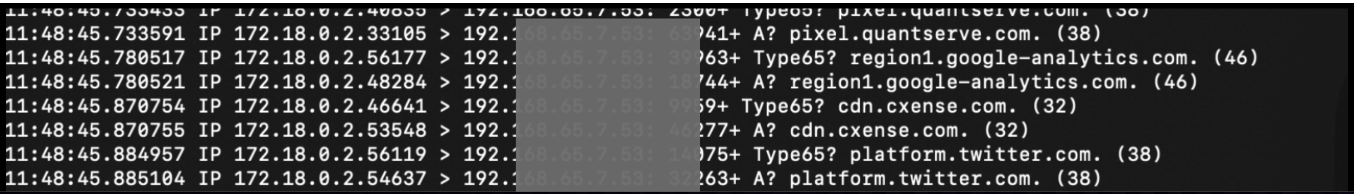

```
11:48:45.733433 IP 172.18.0.2.40835 > 192.168.65.7.53: 2300+ Type65? pixel.quantserve.com. (38)
11:48:45.733591 IP 172.18.0.2.33105 > 192.1[redacted]41+ A? pixel.quantserve.com. (38)
11:48:45.780517 IP 172.18.0.2.56177 > 192.1[redacted]63+ Type65? region1.google-analytics.com. (46)
11:48:45.780521 IP 172.18.0.2.48284 > 192.1[redacted]44+ A? region1.google-analytics.com. (46)
11:48:45.870754 IP 172.18.0.2.46641 > 192.1[redacted]9+ Type65? cdn.cxense.com. (32)
11:48:45.870755 IP 172.18.0.2.53548 > 192.1[redacted]277+ A? cdn.cxense.com. (32)
11:48:45.884957 IP 172.18.0.2.56119 > 192.1[redacted]75+ Type65? platform.twitter.com. (38)
11:48:45.885104 IP 172.18.0.2.54637 > 192.1[redacted]263+ A? platform.twitter.com. (38)
```

OBSERVE THE CROWD
- DO NOT STORE MY BROWSING
to catch the bad guys!

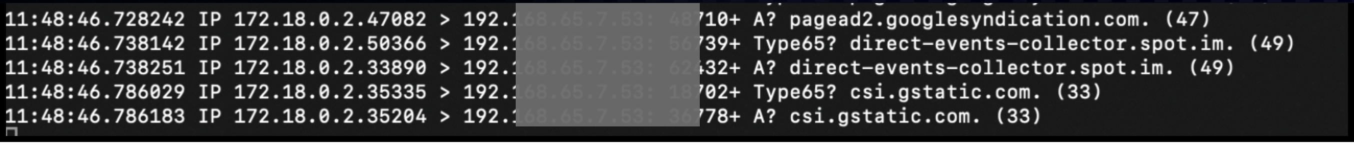

```
11:48:46.728242 IP 172.18.0.2.47082 > 192.1[redacted]10+ A? pagead2.googlesyndication.com. (47)
11:48:46.738142 IP 172.18.0.2.50366 > 192.1[redacted]39+ Type65? direct-events-collector.spot.im. (49)
11:48:46.738251 IP 172.18.0.2.33890 > 192.1[redacted]32+ A? direct-events-collector.spot.im. (49)
11:48:46.786029 IP 172.18.0.2.35335 > 192.1[redacted]02+ Type65? csi.gstatic.com. (33)
11:48:46.786183 IP 172.18.0.2.35204 > 192.1[redacted]78+ A? csi.gstatic.com. (33)
```

Edge Policy Processor

Policy Processor (POP) sets local filtering decisions.

Based on TAPIR observations and other sources

Policy decisions is made by data owner

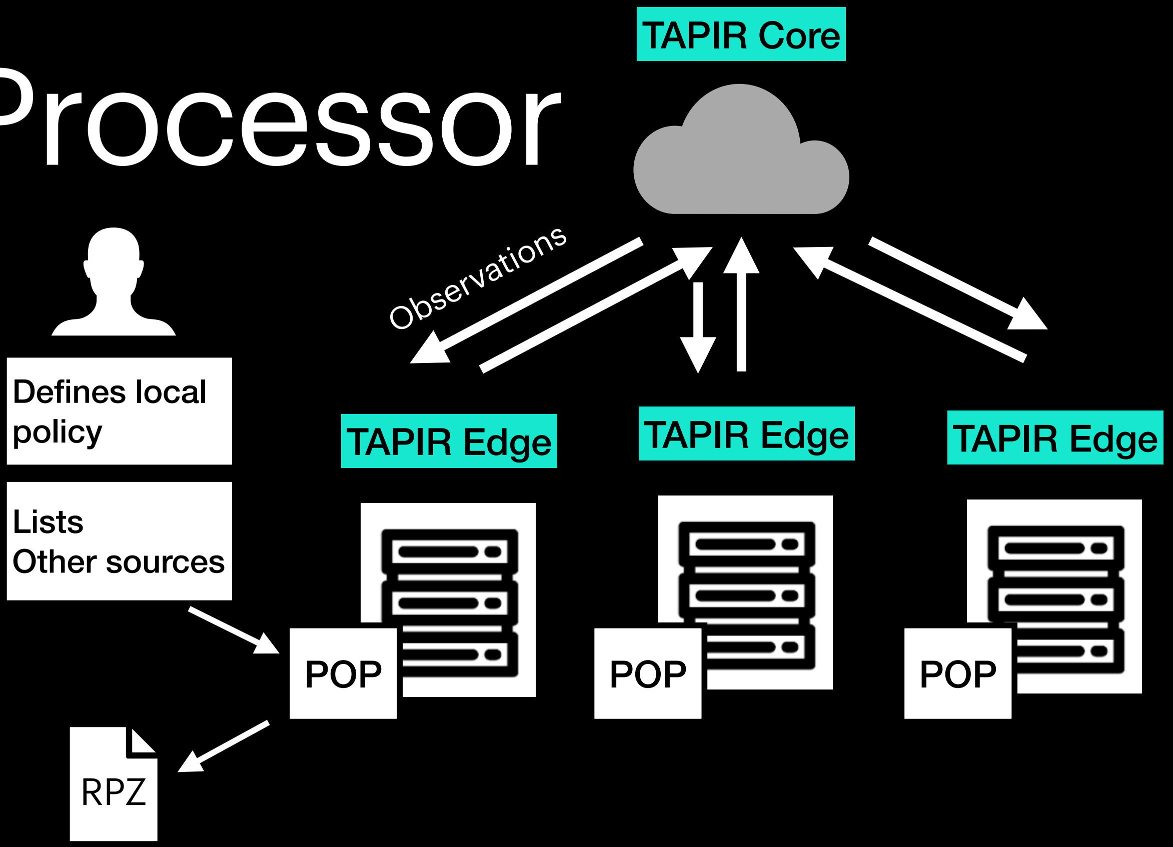

Edge Local Analyse

Analyse of unique domains and sensitive patterns on local resolver. Currently these are discarded.

Tag/categorise domains sent to Core.

Features of new domains saved locally, makes it possible to train a BERT model.

Example:
- changes in botnet C2 domains
- DGAs
- tracking

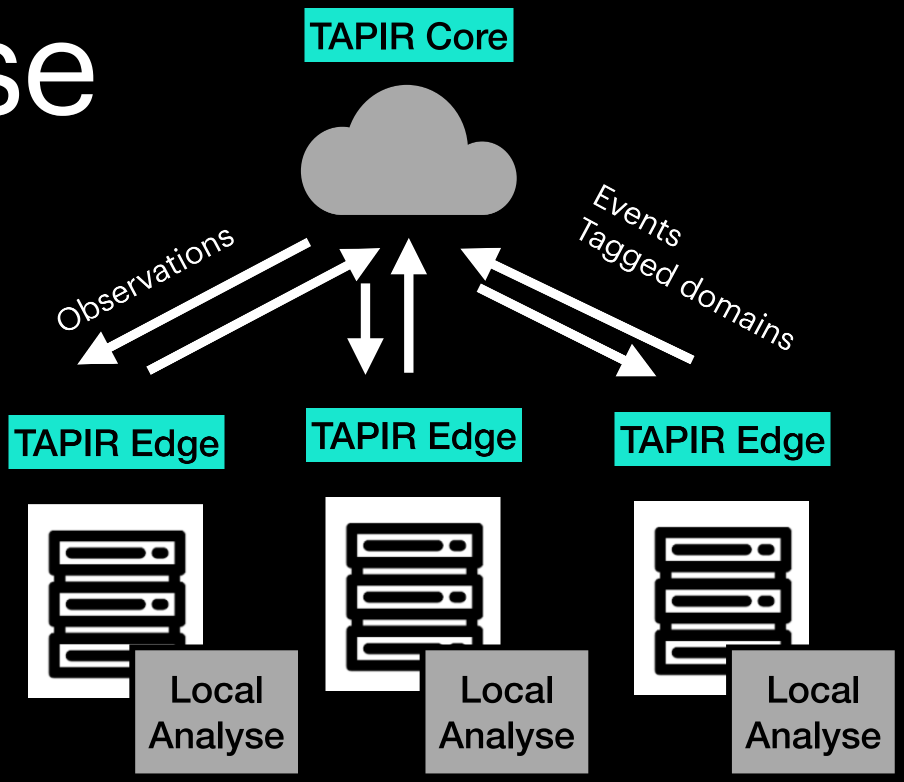

Identifiable browsing behavior never leaves the resolver

”Demo”: The DNS TAPIR loop

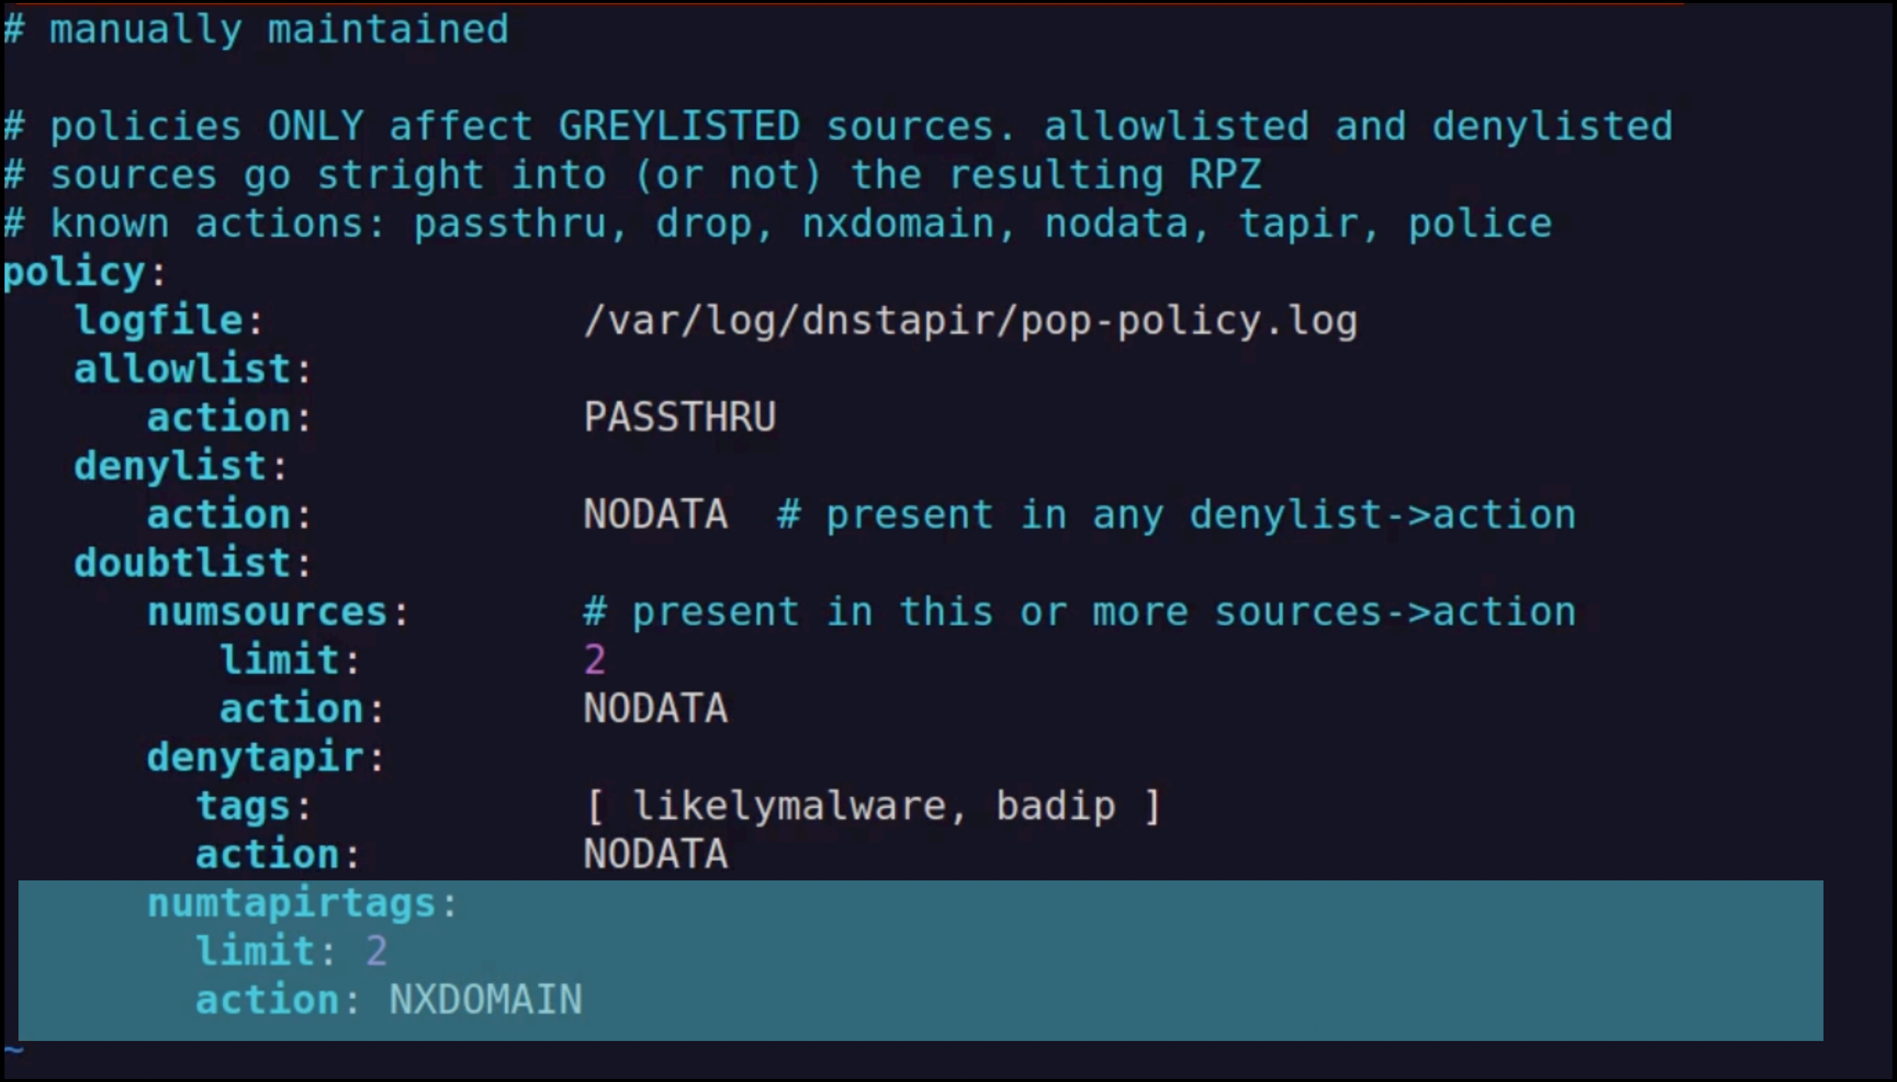

```
# manually maintained

# policies ONLY affect GREYLISTED sources. allowlisted and denylisted
# sources go stright into (or not) the resulting RPZ
# known actions: passthru, drop, nxdomain, nodata, tapir, police
policy:
   logfile:             /var/log/dnstapir/pop-policy.log
   allowlist:
      action:           PASSTHRU
   denylist:
      action:           NODATA  # present in any denylist->action
   doubtlist:
      numsources:       # present in this or more sources->action
         limit:         2
         action:        NODATA
      denytapir:
        tags:           [ likelymalware, badip ]
        action:         NODATA
      numtapirtags:
        limit: 2
        action: NXDOMAIN
~
```

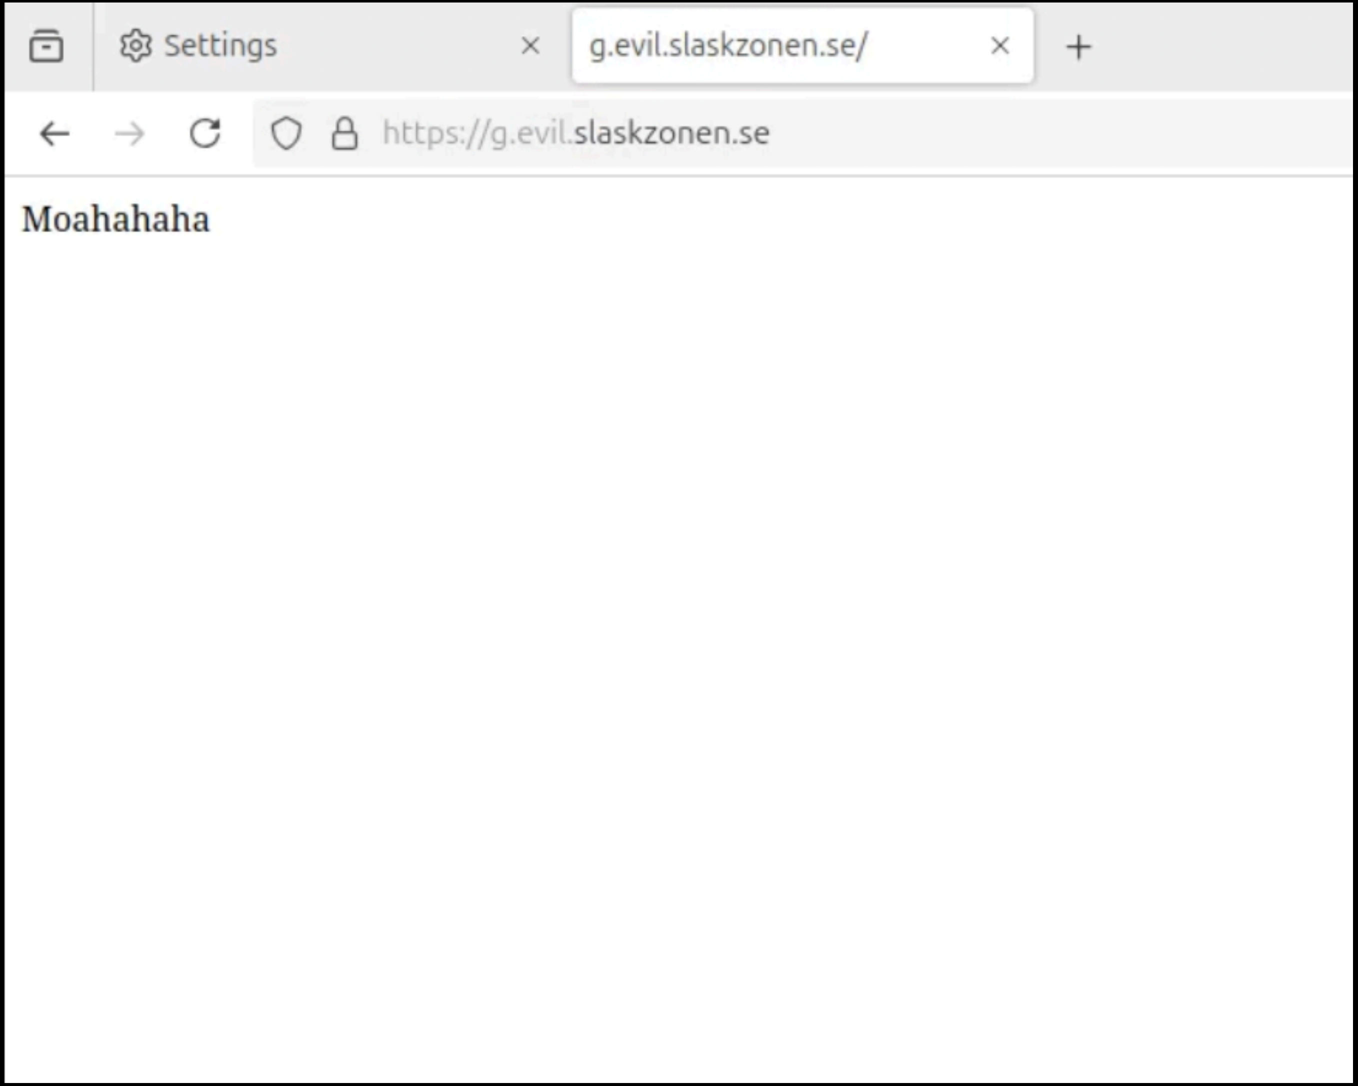

Moahahaha

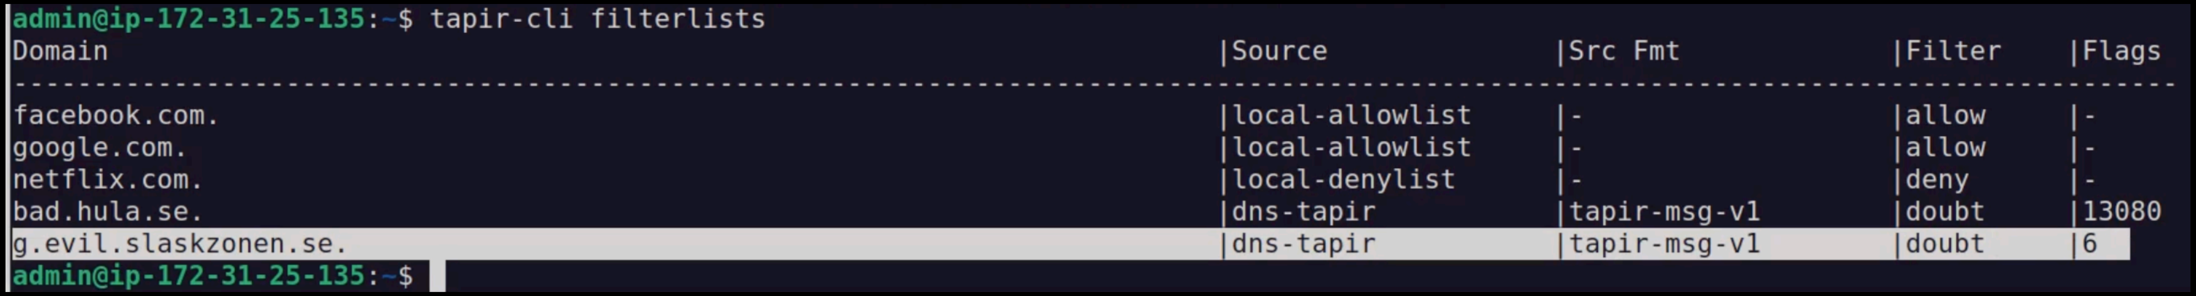

```
admin@ip-172-31-25-135:~$ tapir-cli filterlists
Domain                                                                     |Source              |Src Fmt             |Filter    |Flags
---------------------------------------------------------------------------------------------------------------------------------------
facebook.com.                                                              |local-allowlist     |-                   |allow     |-
google.com.                                                                |local-allowlist     |-                   |allow     |-
netflix.com.                                                               |local-denylist      |-                   |deny      |-
bad.hula.se.                                                               |dns-tapir           |tapir-msg-v1        |doubt     |13080
g.evil.slaskzonen.se.                                                      |dns-tapir           |tapir-msg-v1        |doubt     |6
admin@ip-172-31-25-135:~$
```

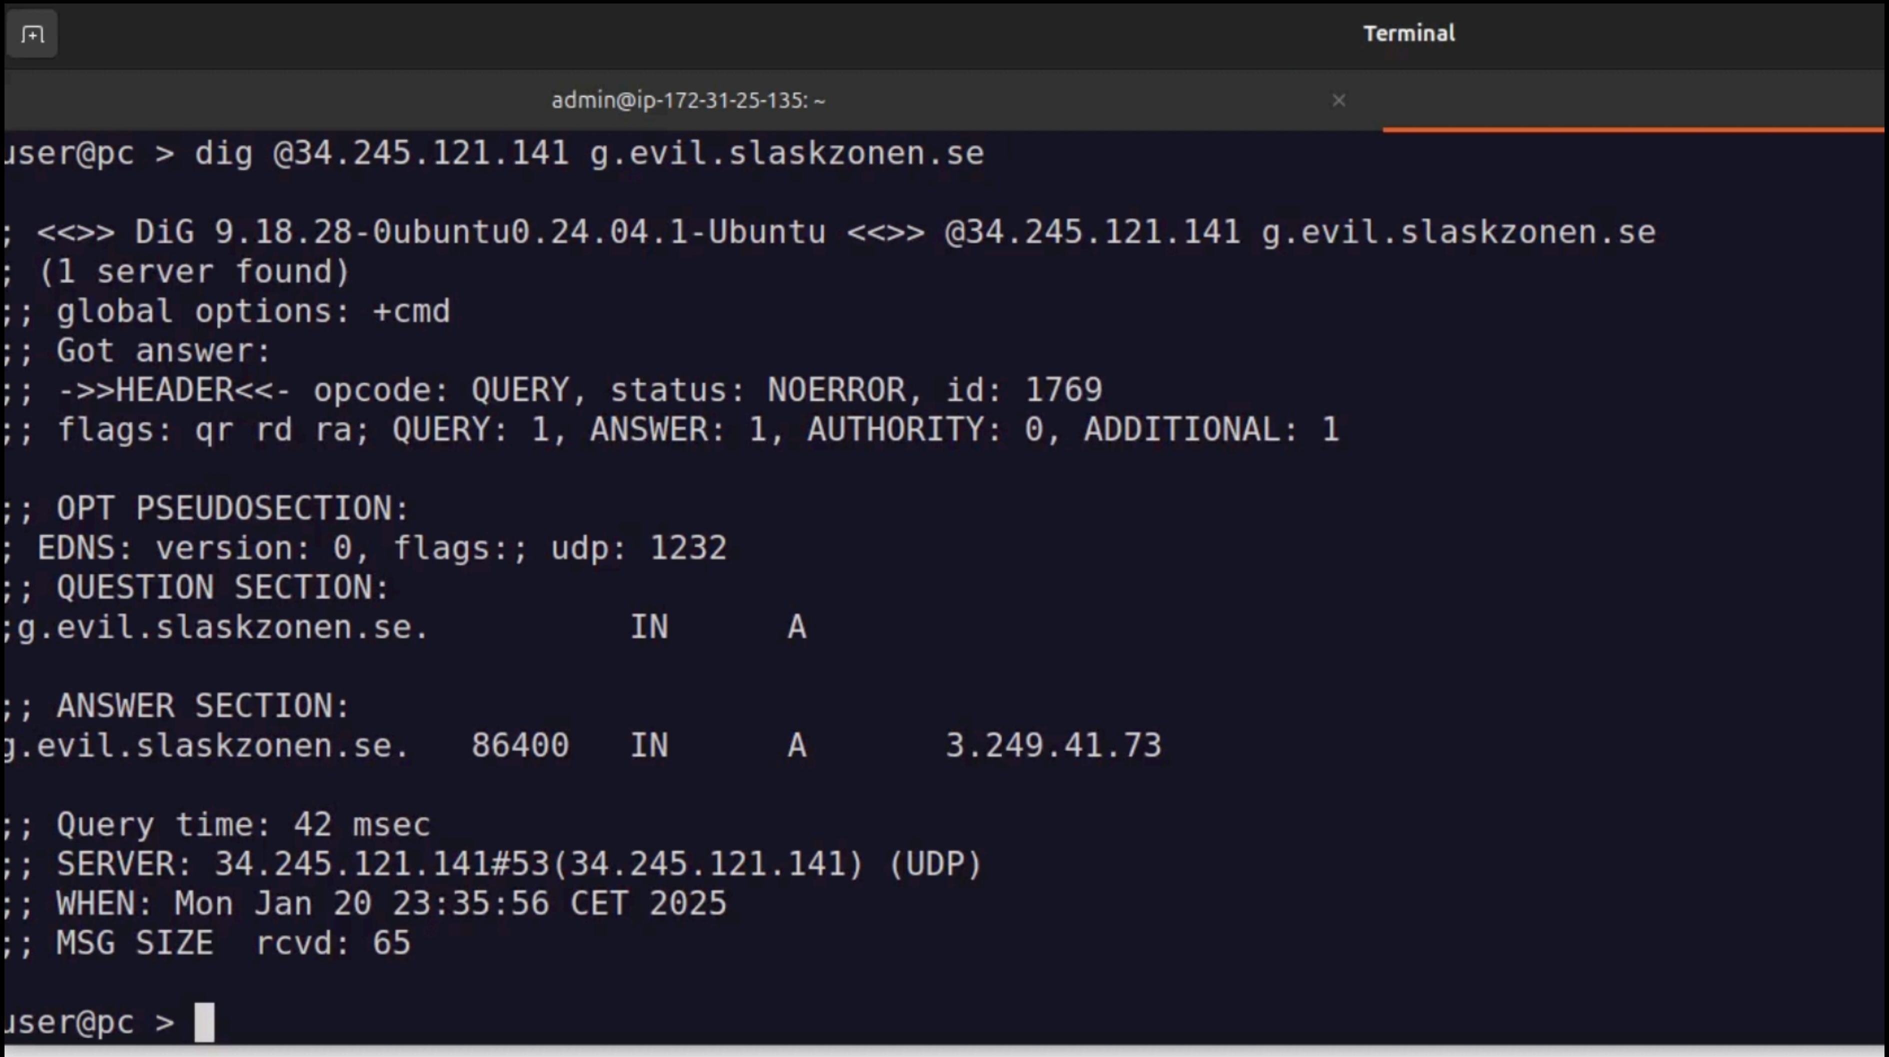

```
ıser@pc > dig @34.245.121.141 g.evil.slaskzonen.se

; <<>> DiG 9.18.28-0ubuntu0.24.04.1-Ubuntu <<>> @34.245.121.141 g.evil.slaskzonen.se
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 1769
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;g.evil.slaskzonen.se.          IN      A

;; ANSWER SECTION:
յ.evil.slaskzonen.se.   86400   IN      A       3.249.41.73

;; Query time: 42 msec
;; SERVER: 34.245.121.141#53(34.245.121.141) (UDP)
;; WHEN: Mon Jan 20 23:35:56 CET 2025
;; MSG SIZE  rcvd: 65

ıser@pc > 
```

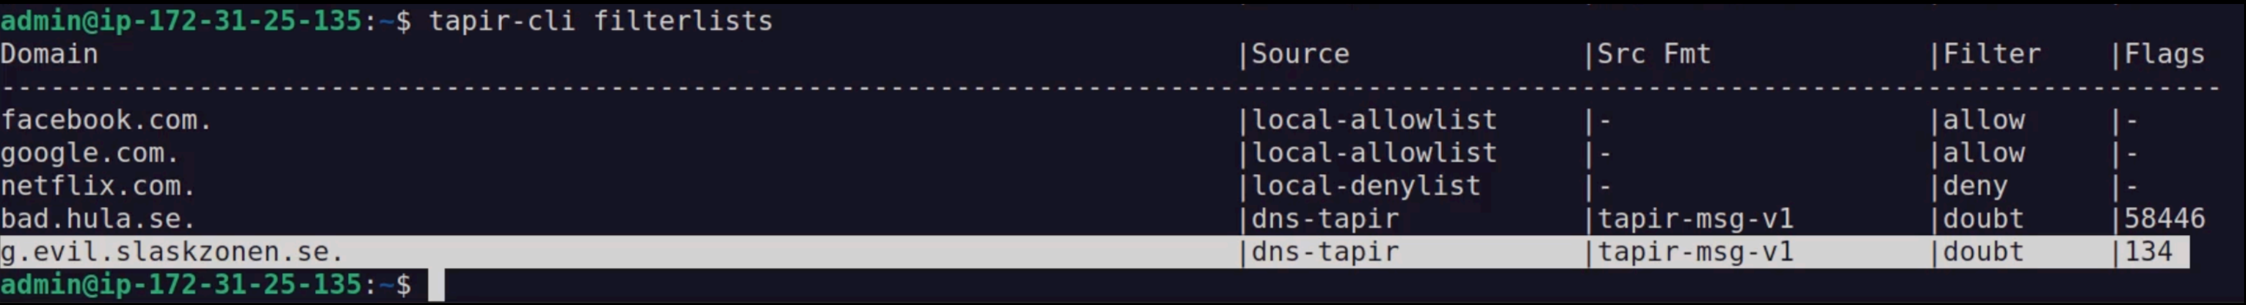

```
admin@ip-172-31-25-135:~$ tapir-cli filterlists
Domain                                                                     |Source              |Src Fmt             |Filter    |Flags
---------------------------------------------------------------------------------------------------------------------------------------
facebook.com.                                                              |local-allowlist     |-                   |allow     |-
google.com.                                                                |local-allowlist     |-                   |allow     |-
netflix.com.                                                               |local-denylist      |-                   |deny      |-
bad.hula.se.                                                               |dns-tapir           |tapir-msg-v1        |doubt     |58446
g.evil.slaskzonen.se.                                                      |dns-tapir           |tapir-msg-v1        |doubt     |134
admin@ip-172-31-25-135:~$
```

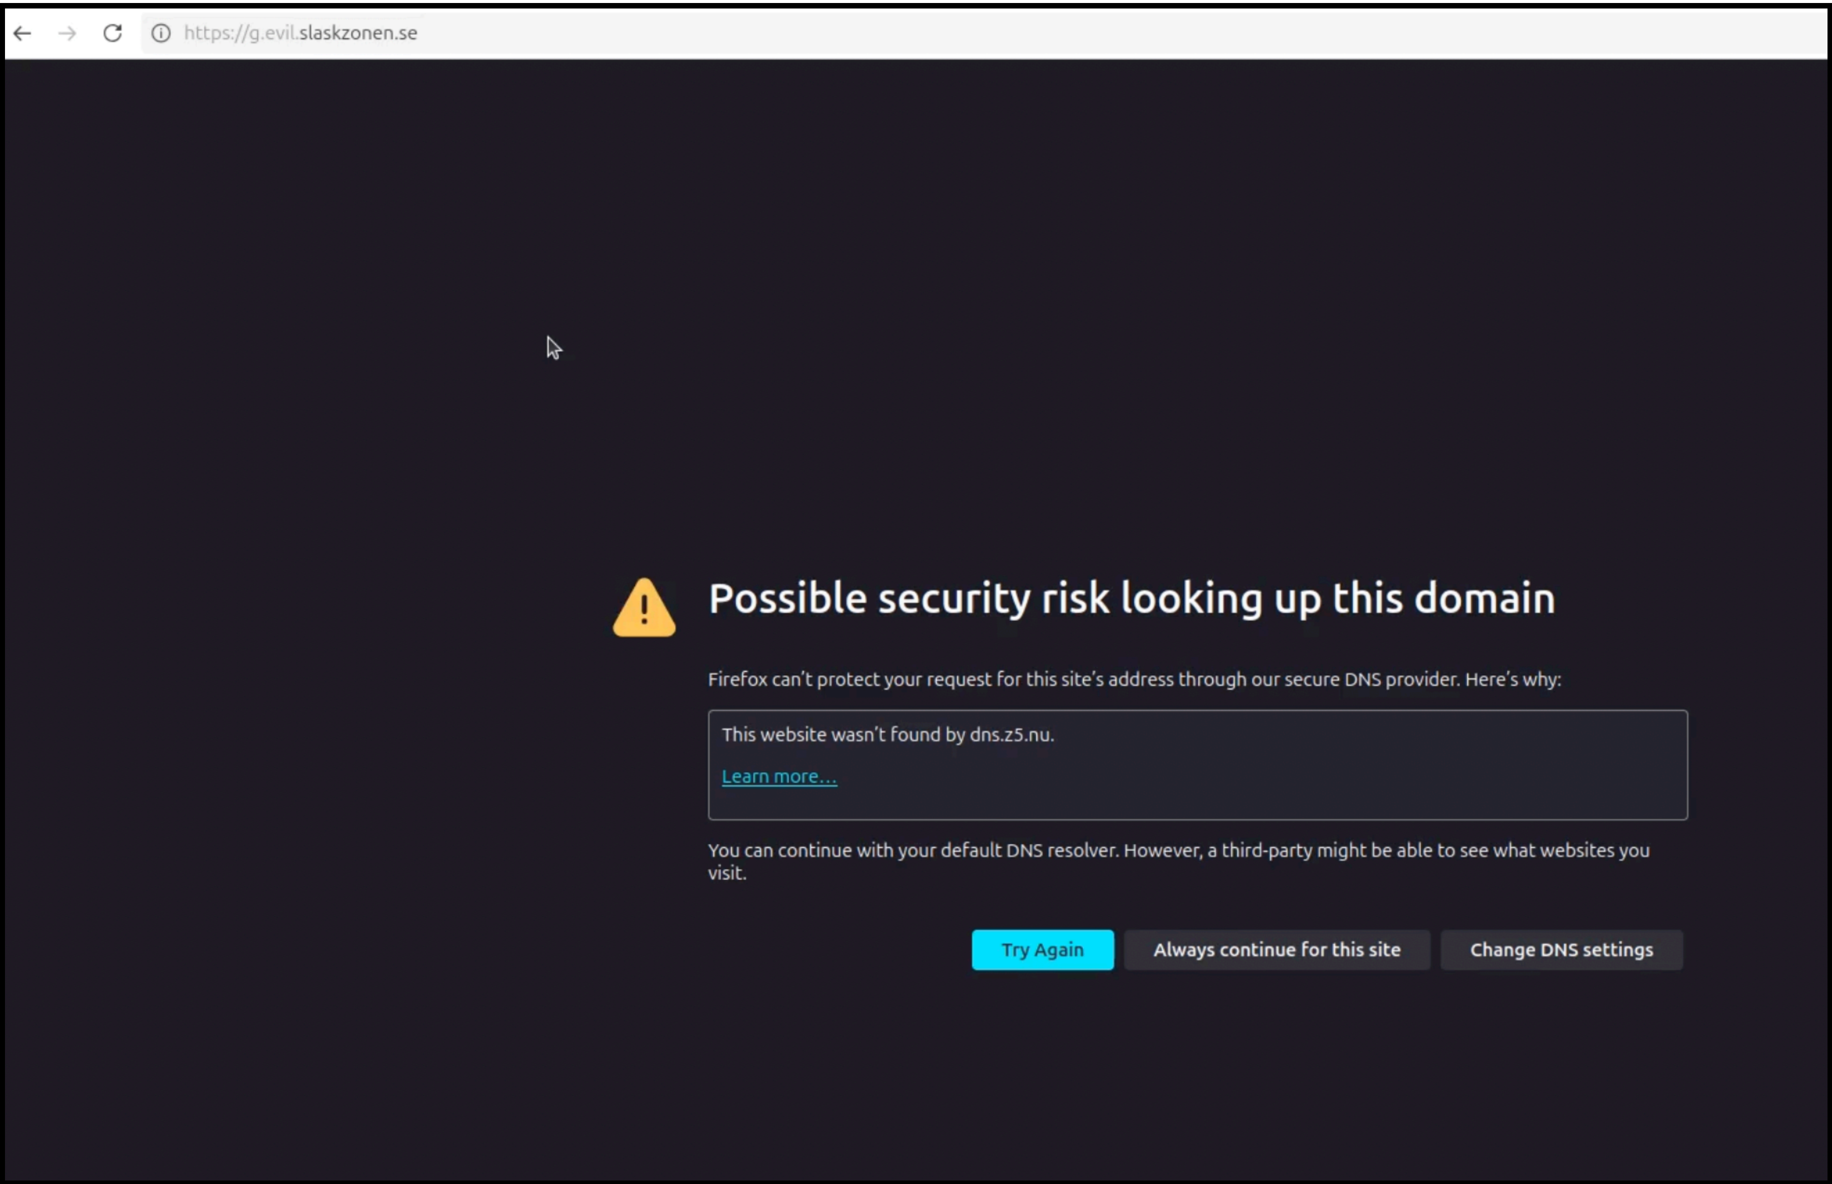

Possible security risk looking up this domain

Firefox can't protect your request for this site's address through our secure DNS provider. Here's why:

This website wasn't found by dns.z5.nu.

Learn more…

You can continue with your default DNS resolver. However, a third-party might be able to see what websites you visit.

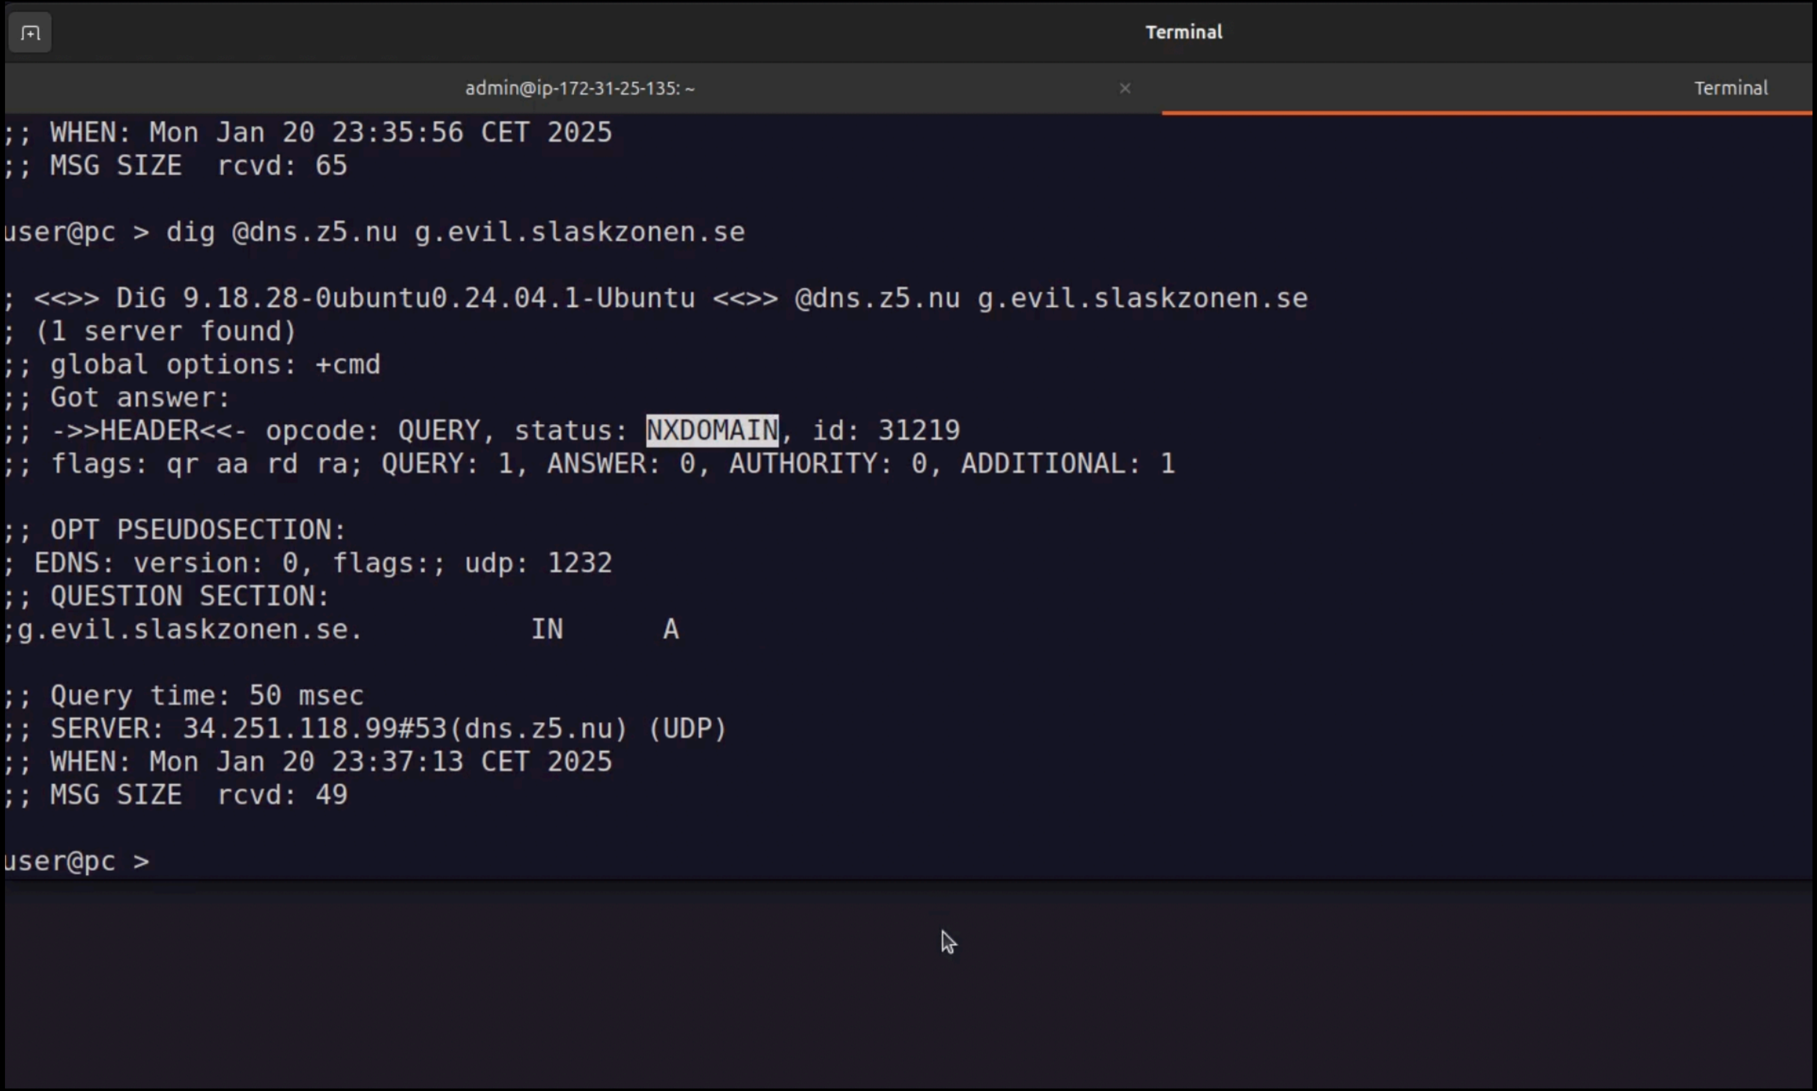

```
;; WHEN: Mon Jan 20 23:35:56 CET 2025
;; MSG SIZE  rcvd: 65

user@pc > dig @dns.z5.nu g.evil.slaskzonen.se

; <<>> DiG 9.18.28-0ubuntu0.24.04.1-Ubuntu <<>> @dns.z5.nu g.evil.slaskzonen.se
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 31219
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;g.evil.slaskzonen.se.          IN      A

;; Query time: 50 msec
;; SERVER: 34.251.118.99#53(dns.z5.nu) (UDP)
;; WHEN: Mon Jan 20 23:37:13 CET 2025
;; MSG SIZE  rcvd: 49

user@pc >
```

MAKE INSTALL?

info@dnstapir.se

GitHub.com/dnstapir

Building a TAPIR Community

DNS TAPIR
Community chat

https://www.dnstapir.se

dnstapir.se LinkedIn @dnstapir@mastodon.social

Be careful
what you ask for

DNS TAPIR
WWW.DNSTAPIR.SE

ulrika.vincent@agical.se

Questions?

Images by:

https://www.hive.co.uk/Product/The-Good-the-Bad-and-the-Ugly/32122031

macrovector on Freepik

https://manwithnoname.fandom.com/wiki/Tuco

https://farm5.staticflickr.com/4693/39372851911_72ef64558b.jpg

https://webhostinggeeks.com/guides/privacy/

https://bamfstyle.com/2020/05/31/good-bad-ugly-clint-eastwood/

https://cansupeker.com/the-good-the-bad-and-the-ugly-1966/

https://www.thevintagenews.com/2018/12/20/lee-van-cleef/