Akamai

DNS Security Analytics

and/or Privacy

Mikael Kullberg

ME

- ★ Internet Jack-of-all-trades
- ★ Currently enacting a Security Research/Data Science/Data Engineering hybrid role at Akamai Technologies
- ★ DNS-related threat research since 2014
- ★ Likes open-ended datasets, minimal ground truth and dynamic inputs

Agenda

DNS filtering

DNS attacks on DNS, and Botnets
Malware and phishing

DNS log analytics

Collection of data
Processing of data

DNS privacy

Pseudonymization, retention, aggregation
Mosaic effect, client-specific queries

DNS filtering: DNS attacks and botnets

Attacks on DNS with DNS

- DNS Amplification attacks / Reflection
- PRSD / Chinese water torture

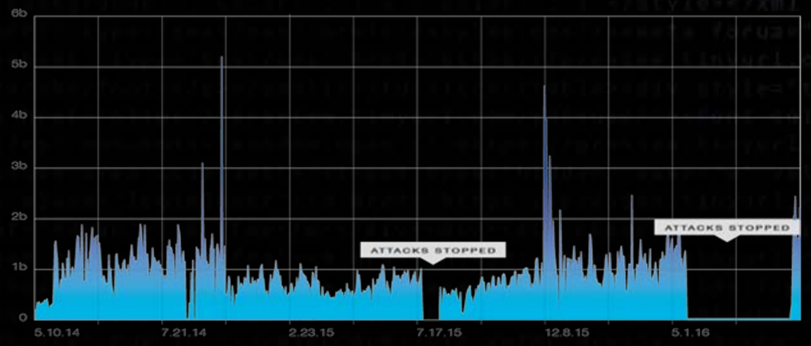

Botnet C2 traffic

- ISP Clients are (not all, but frequently) bots
- Bots call home for instructions
- No instructions, no bot activity
- Profit!

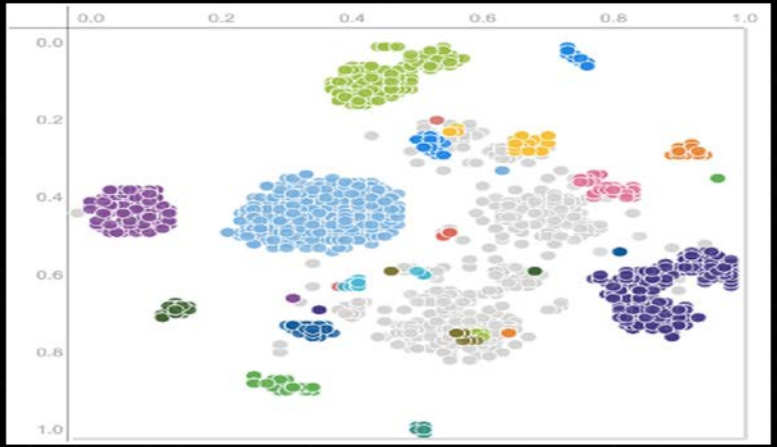

From the ISP perspective, this is network hygiene

DNS filtering: Malware and phishing

Bot or malware? A bot is malware and (almost) everything is a RAT. Not all RATs are bots. Not all bots are RATs. All are bad, all are software, hence malware.

So what gives?

Classification is hard…

Two criteria: Subscriber perspective and pre-infection

Malware

- Links in emails, text messages, etc, also links on webpages, malvertizing and similar
- Drive-by downloads, JS malware, malware repositories

Phishing

- Links in emails, text messages, chat, etc actively being pushed to the user
- Links to sites pretending to be something else to steal your credentials

Classification, it seems, is also somewhat arbitrary…

DNS filtering: Don’t be late

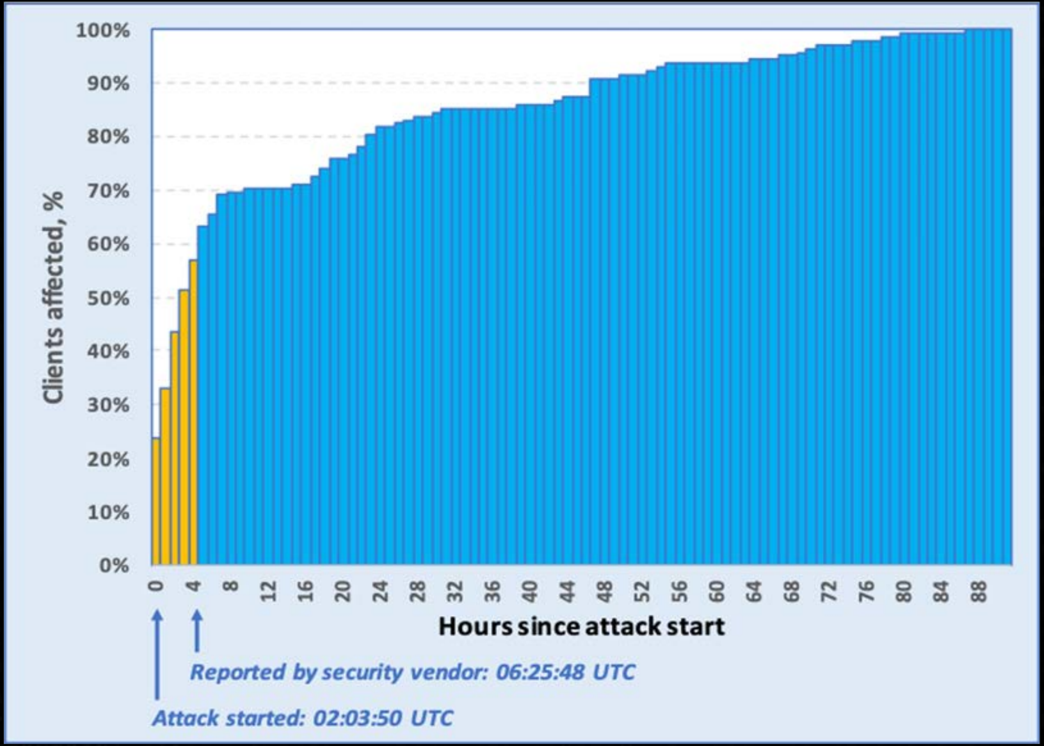

- **Attacks are getting shorter**
- **Attacks give returns early**
- **Attacks are regional**

The chart shows a spray-and-pray phishing attack against a regional bank for clients limited to that region, and this attack is considered slow today. Old data, shown for basic premise.

For example, the more recent **Flubot** malware generates phishing links over text messages, and the DNS records are only actively pointing at malware servers for as little as **10 minutes**.

How does that work with Precision versus Recall?

DNS filtering: One of our petri dishes…

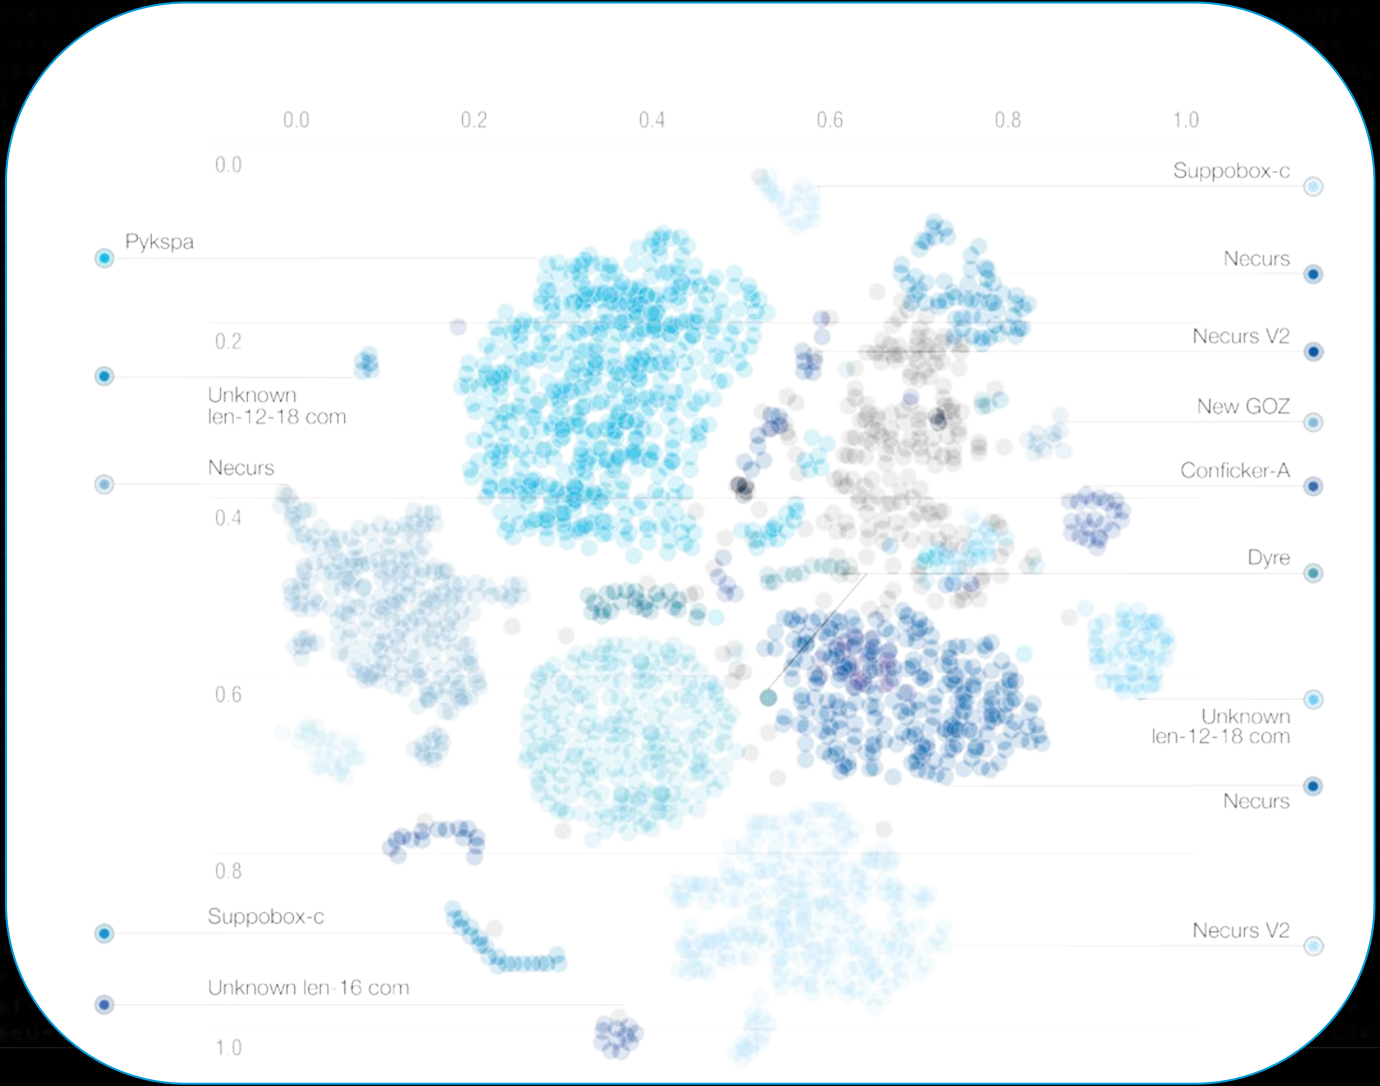

DNS analytics: collecting

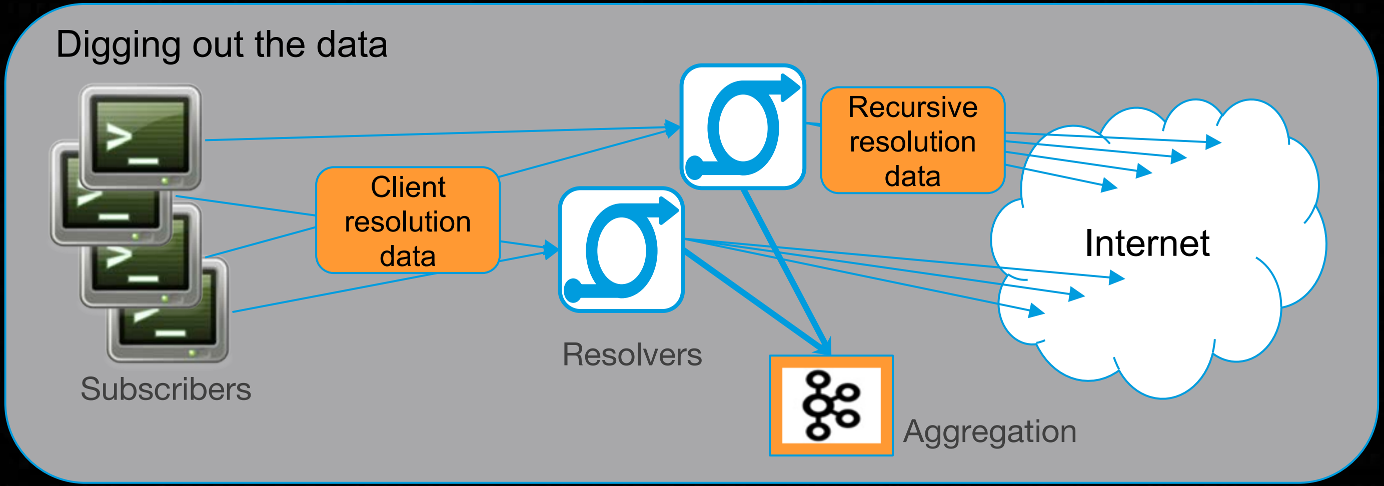

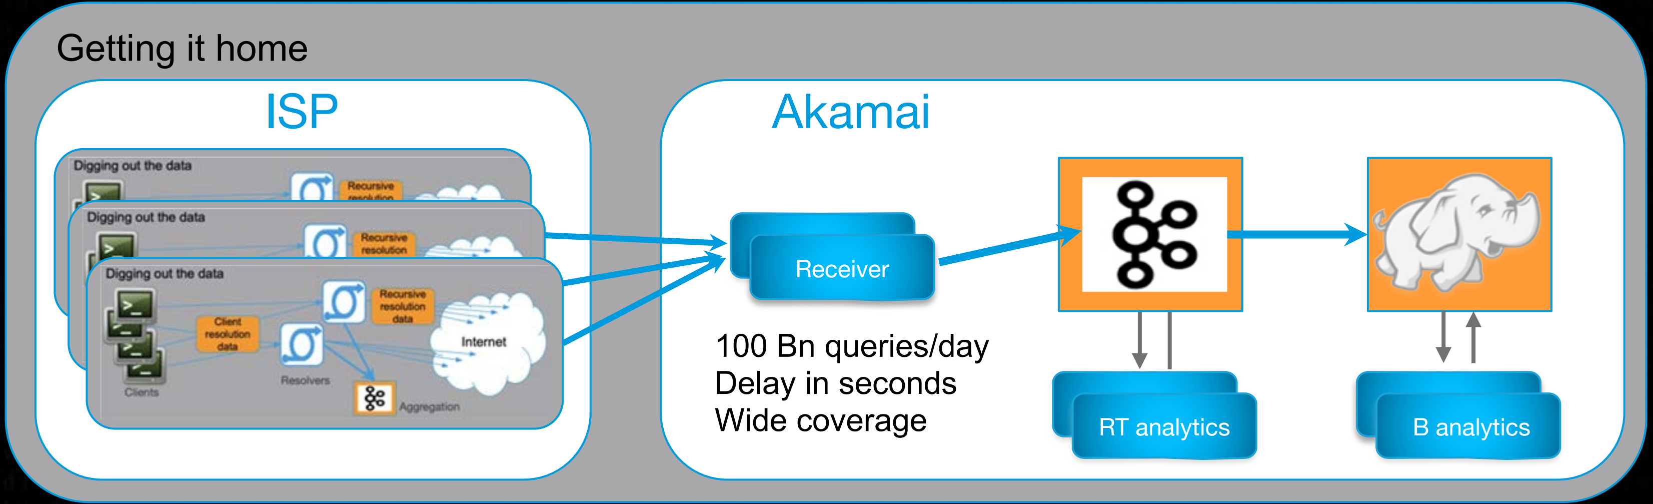

100 Bn queries/day
Delay in seconds
Wide coverage

DNS analytics: processing

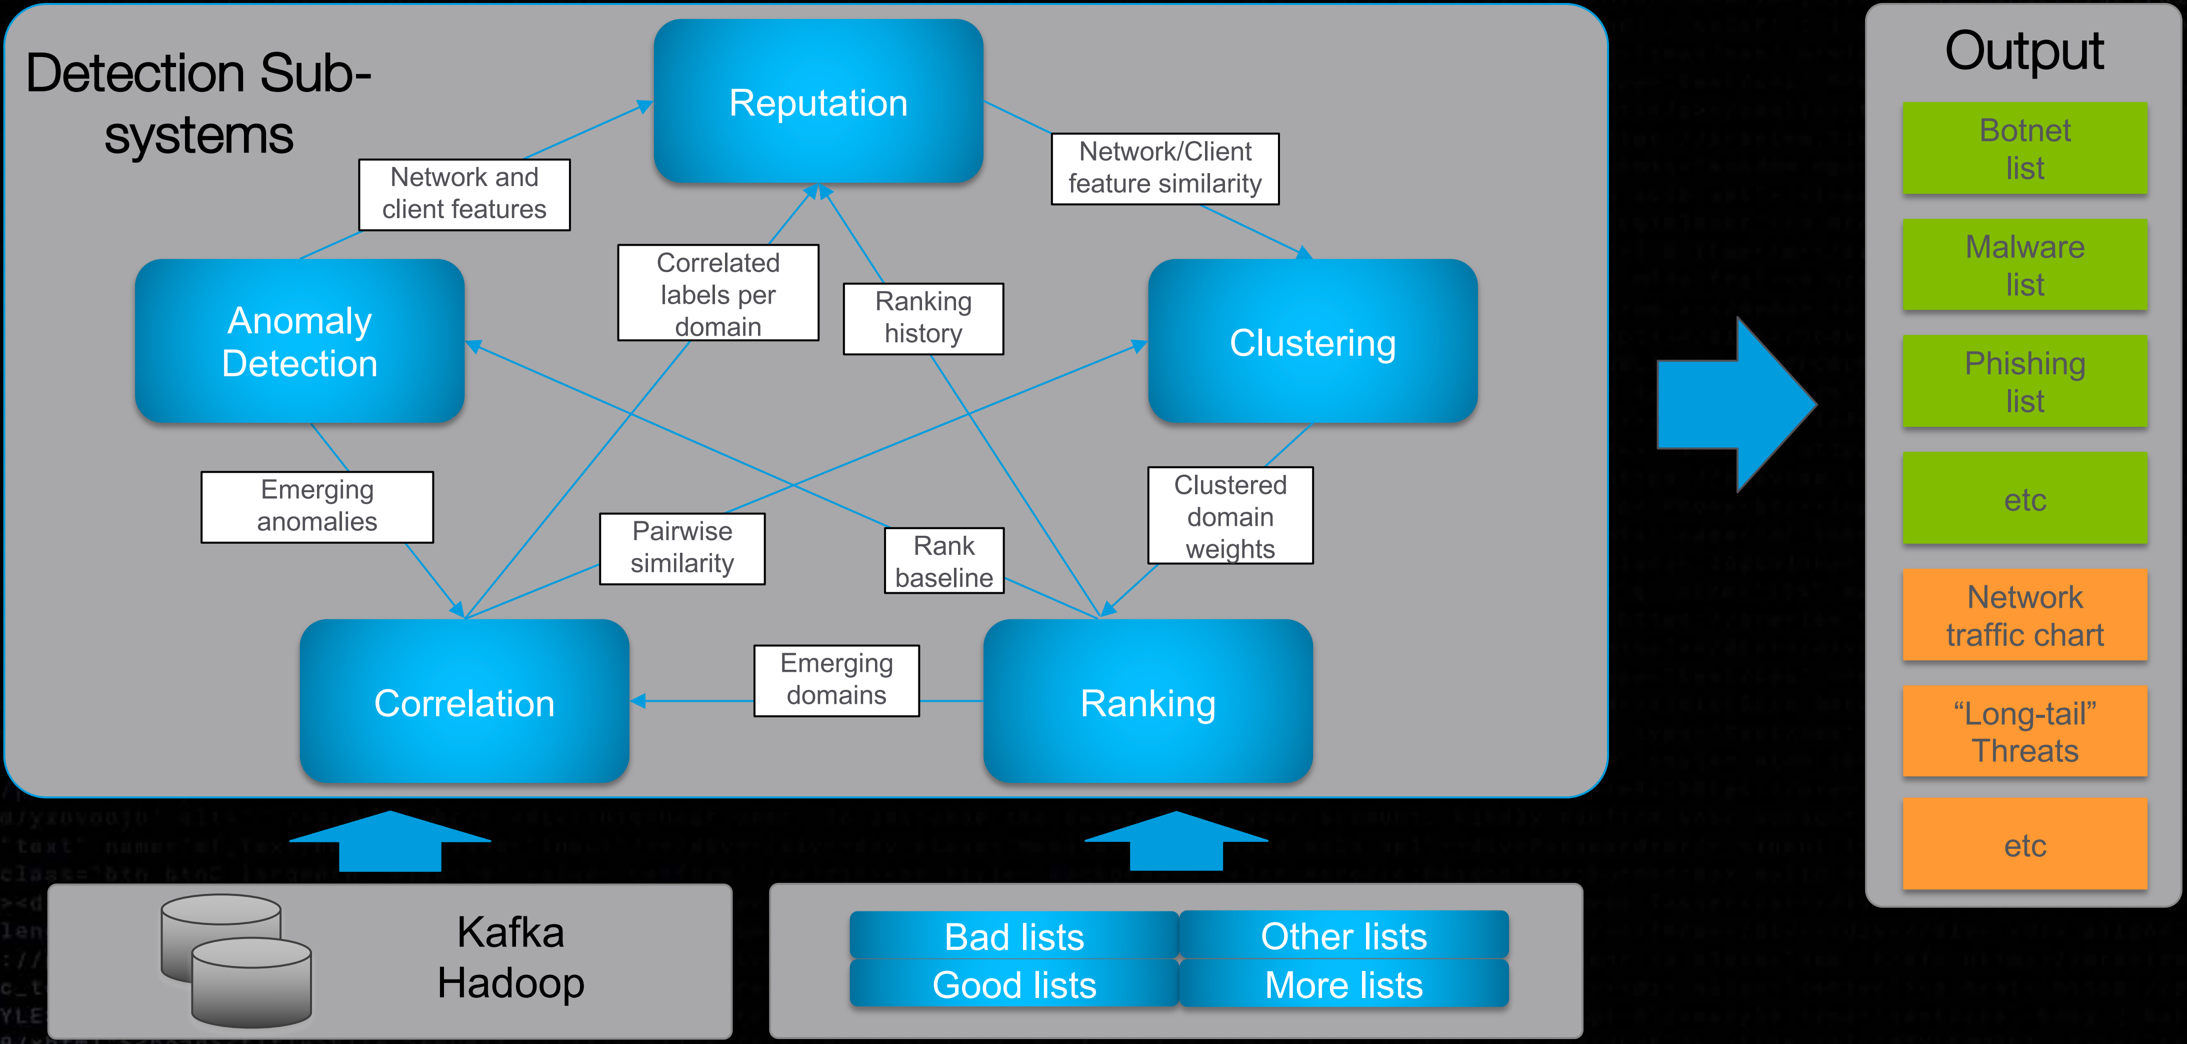

DNS Privacy: what we do

Pseudonymization

Raw data

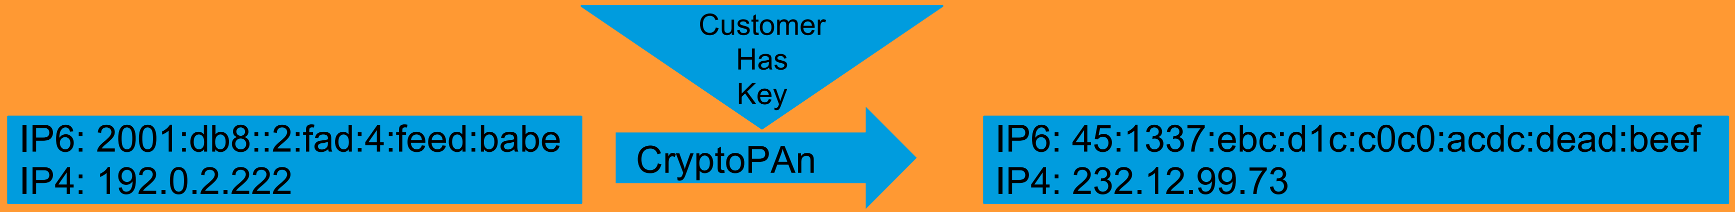

IP6: 2001:db8::2:fad:4:feed:babe
IP4: 192.0.2.222

IP6: 45:1337:ebc:d1c:c0c0:acdc:dead:beef
IP4: 232.12.99.73

This is baseline. Some carriers go further, with higher key rotation frequency and lumping subscriber IPv4s within a /28 together.

Intermediary data

For analytics processes that retain a query - client relationship, use CryptoPAn again to dissociate the key from raw data

Retention

Raw data

What you don’t have, you can’t lose. Raw logs with client pseudonyms are processed, aggregated and discarded within a 7 day window.

Intermediary data

Multiple queries connected to a client pseudonym must not retain query order and/or precise timestamps and are kept for 30 days.

Aggregation

Raw data

Feature extraction per client is only seen as relevant for 24h. Client pseudonym is not a reliable grouping key beyond that window due to DHCP, CryptoPAn key rotation, CGNAT, etc. Beyond that, aggregation is statistical per FQDN where clients are reduced to a quantity feature.

Intermediary data

Baselines based on client pseudonyms are meaningless as persistence is unreliable and random. All long term data is related to domain names only.

DNS Privacy: Is it enough?

For our customers who share data with us? Yes (obviously)
For GDPR compliance? Yes

For you? Depends on your threat model.

The Mosaic effect

If there are multiple logs related to a subject that are loosely connected by time, events, etc, then cross-referencing will yield additional information and potentially break pseudonymization

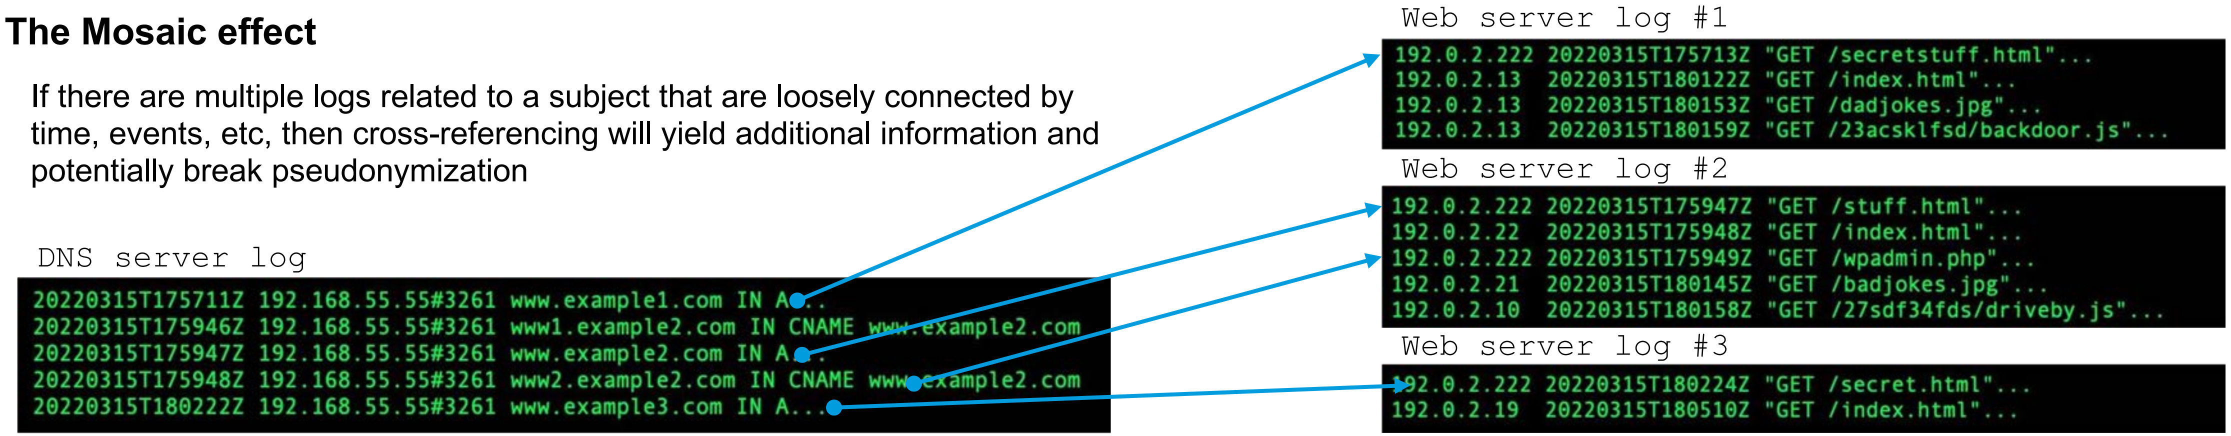

DNS server log

```
20220315T175711Z 192.168.55.55#3261 www.example1.com IN A..
20220315T175946Z 192.168.55.55#3261 www1.example2.com IN CNAME www.example2.com
20220315T175947Z 192.168.55.55#3261 www.example2.com IN A..
20220315T175948Z 192.168.55.55#3261 www2.example2.com IN CNAME www.example2.com
20220315T180222Z 192.168.55.55#3261 www.example3.com IN A...
```

Web server log #1

```
192.0.2.222 20220315T175713Z "GET /secretstuff.html"...
192.0.2.13  20220315T180122Z "GET /index.html"...
192.0.2.13  20220315T180153Z "GET /dadjokes.jpg"...
192.0.2.13  20220315T180159Z "GET /23acsklfsd/backdoor.js"...
```

Web server log #2

```
192.0.2.222 20220315T175947Z "GET /stuff.html"...
192.0.2.22  20220315T175948Z "GET /index.html"...
192.0.2.222 20220315T175949Z "GET /wpadmin.php"...
192.0.2.21  20220315T180145Z "GET /badjokes.jpg"...
192.0.2.10  20220315T180158Z "GET /27sdf34fds/driveby.js"...
```

Web server log #3

```
192.0.2.222 20220315T180224Z "GET /secret.html"...
192.0.2.19  20220315T180510Z "GET /index.html"...
```

Individualized queries

Subdomains with long, cryptic labels can be any number of things; signatures, dns tunnels, PRSD attacks, etc. The examples are seen repeatedly from single clients. Queries have been garbled to protect both the client and the guilty.

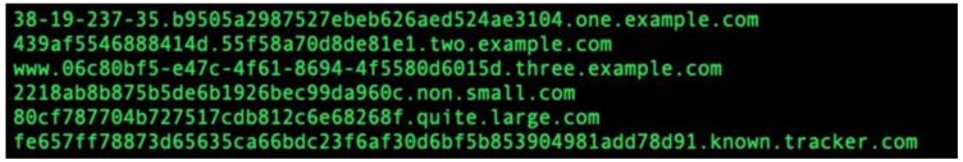

```
38-19-237-35.b9505a2987527ebeb626aed524ae3104.one.example.com
439af5546888414d.55f58a70d8de81e1.two.example.com
www.06c80bf5-e47c-4f61-8694-4f5580d6015d.three.example.com
2218ab8b875b5de6b1926bec99da960c.non.small.com
80cf787704b727517cdb812c6e68268f.quite.large.com
fe657ff78873d65635ca66bdc23f6af30d6bf5b853904981add78d91.known.tracker.com
```

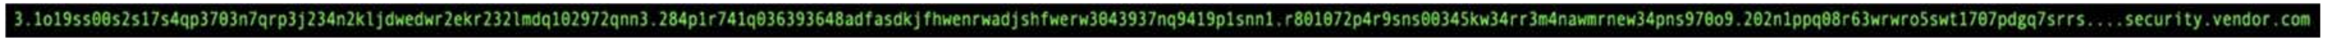

```
3.1o19ss00s2s17s4qp3703n7qrp3j234n2kljdwedwr2ekr232lmdq102972qnn3.284p1r741q036393648adfasdkjfhwenrwadjshfwerw3043937nq9419p1snn1.r801072p4r9sns00345kw34rr3m4nawmrnew34pns970o9.202n1ppq08r63wrwro5swt1707pdgq7srrs....security.vendor.com
```

Summary

- DNS filtering is most effective for phishing, pre-infection malware and post-infection information leakage/C2. Still a relevant security control.
- Threat research aims for maximum utility – difficult knife's edge between type I and type II errors. False positives annoy ISPs, false negatives annoy subscribers.
- Grouping key is irrelevant, it's the grouping that matters.
- Mosaic effect makes CDNs really scary, but any filtering DNS resolver provider can generate its own fingerprints.
- Weird-looking subdomains in queries can upload many bits of data and responses need not be what they seem to be.

Questions?