In rDNS We Trust

Revisiting rDNS Use by Clients on the Internet

Tobias Fiebig[1,2,3], Kevin Borgolte[2], Shuang Hao[4], Christopher Kruegel[2], Giovanni Vigna[2], Anja Feldmann[3,5]

[1] TU Delft
[2] UC Santa Barbara
[3] TU Berlin
[4] UT Dallas
[5] Max Planck Institute for Informatics

TU Delft

What is reverse DNS?

Forward DNS

- Usually gives you an IPv(4|6) Address for a name (A|AAAA query):

Reverse DNS (rDNS)

- Gives you a pointer to a name in the tree:

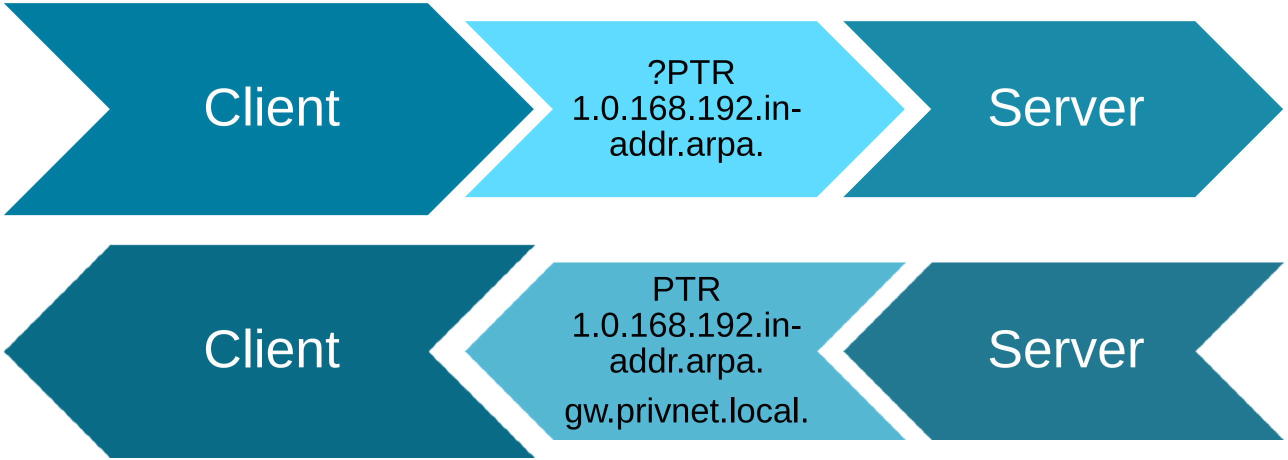

rDNS Zones

- IPv4:
  - in-addr.arpa.
  - Four octets
  - One level per octet:

    xxx.xxx.xxx.xxx.in-addr.arpa.

- IPv6:
  - ip6.arpa.
  - 32 nibbles
  - One level per nibble:

    x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.x.ip6.arpa.

TU Delft


WHATEVER
WHO CARES

You can use it to understand links…

…topologies…

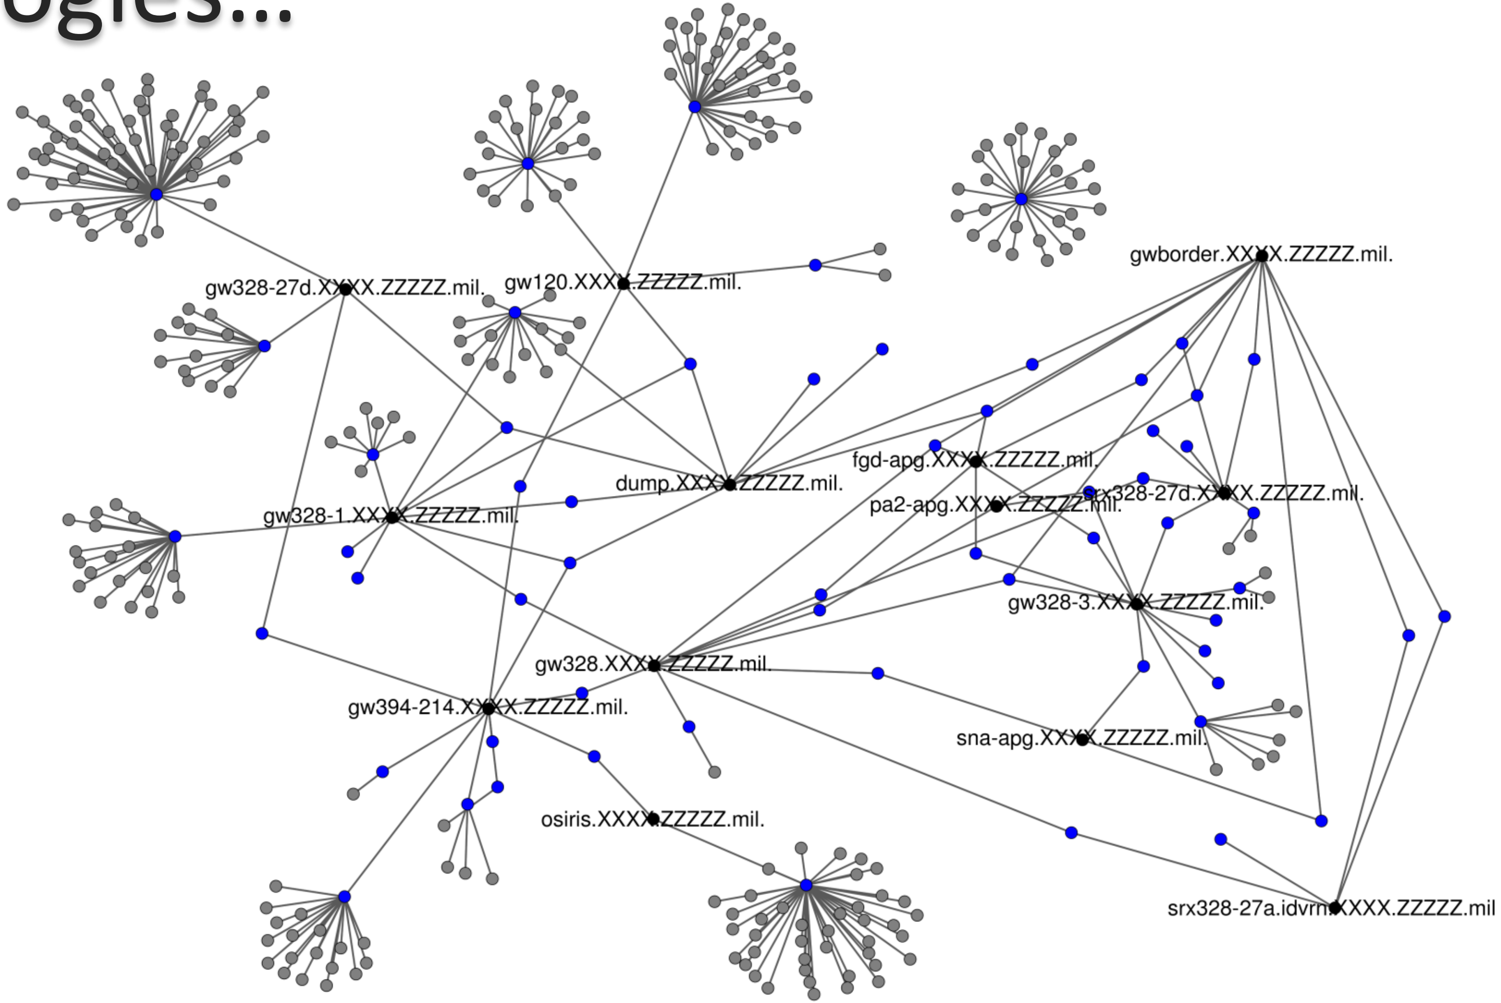

TU Delft

…IPv6 deployments…

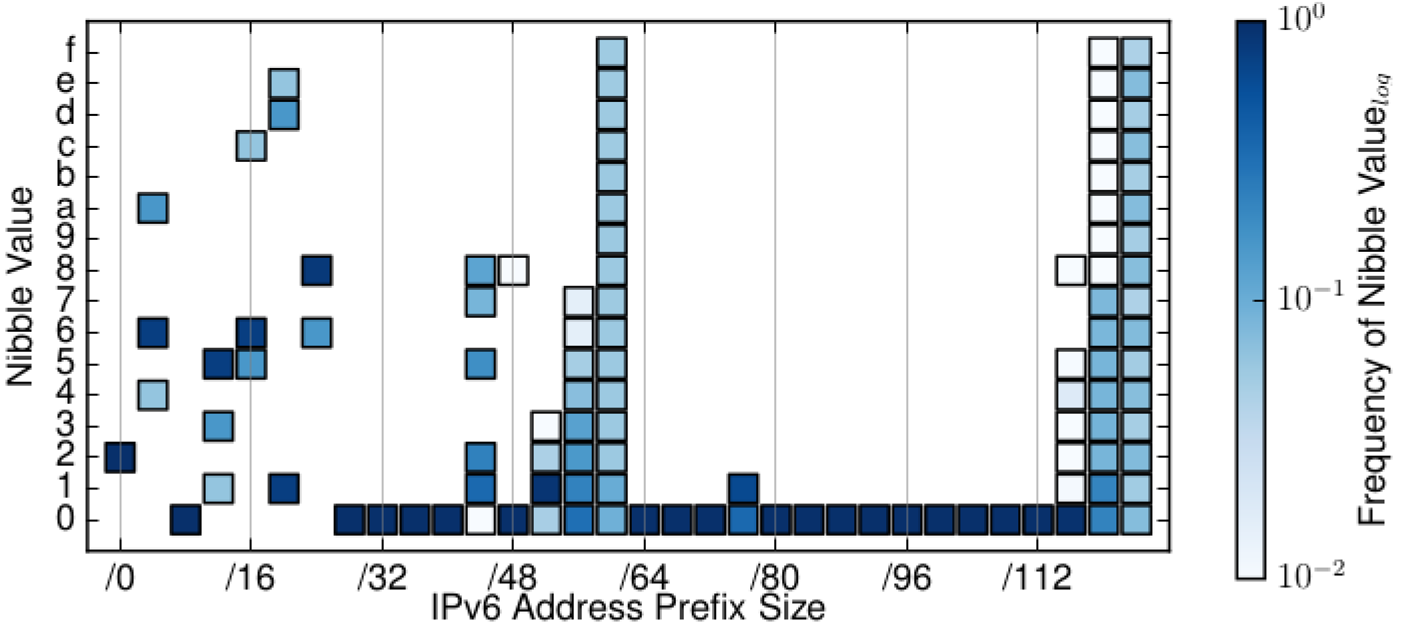

TUDelft

…even build IPv6 security scan seeds

```
tfiebig@shells ~ %
tfiebig@shells ~ % nmap -6 -n -A -T insane 2a01:4f8:10b:37ef::186


Starting Nmap 7.60 ( https://nmap.org ) at 2018-03-26 14:03 CEST
Nmap scan report for 2a01:4f8:10b:37ef::186
Host is up (0.00056s latency).
Not shown: 992 closed ports
PORT      STATE SERVICE   VERSION
22/tcp    open  ssh       OpenSSH 7.6 (protocol 2.0)
| ssh-hostkey:
|   2048 2e:2b:9a:8c:a8:4f:8e:ad:4c:b4:b5:cf:7b:ba:37:8b (RSA)
|   256 fb:8a:5f:0e:d4:6c:a6:c8:45:31:1a:e1:a1:1c:34:5f (ECDSA)
|_  256 5b:a6:f8:e6:d2:4b:c5:c5:d4:64:78:19:d8:44:7a:92 (EdDSA)
25/tcp    open  smtp      Postfix smtpd
|_smtp-commands: mail.aperture-labs.org, PIPELINING, SIZE 20480000, VRFY, ETRN, STARTTLS, ENH
ANCEDSTATUSCODES, 8BITMIME, DSN,
| ssl-cert: Subject: commonName=mail.aperture-labs.org
| Subject Alternative Name: DNS:mail.aperture-labs.org
| Not valid before: 2018-02-26T00:31:56
|_Not valid after:  2018-05-27T00:31:56
|_ssl-date: TLS randomness does not represent time
```

TU Delft

How does real-world rDNS use look?

- Earlier work, e.g., Gao et al. suggests no
  - Argument: High SERVFAIL share for PTR requests indicates low maintenance state

- Still:
  - Not focusing on rDNS
  - More an afterthought of “real” DNS

TUDelft

Passive trace results

Data source: Farsight DNS stream

- Collected from DNS recursors around the globe
- Provided to researchers and IT security professionals
- Large ☺

TUDelft

Passive DNS Data Flow

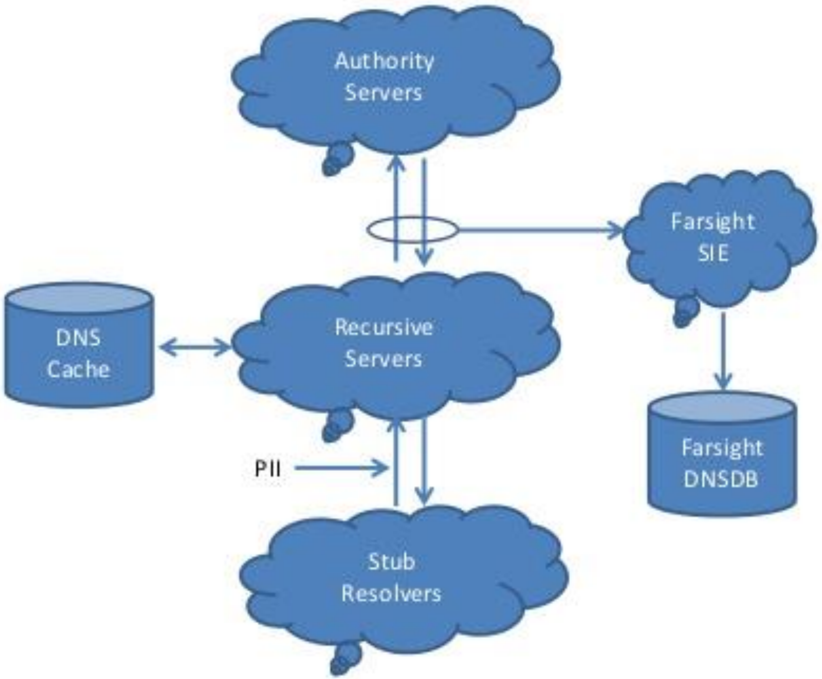

(Slide referenced from: https://www.slideshare.net/apnic/2014-03-dnstap-1410824032)

TU Delft

Full Dataset

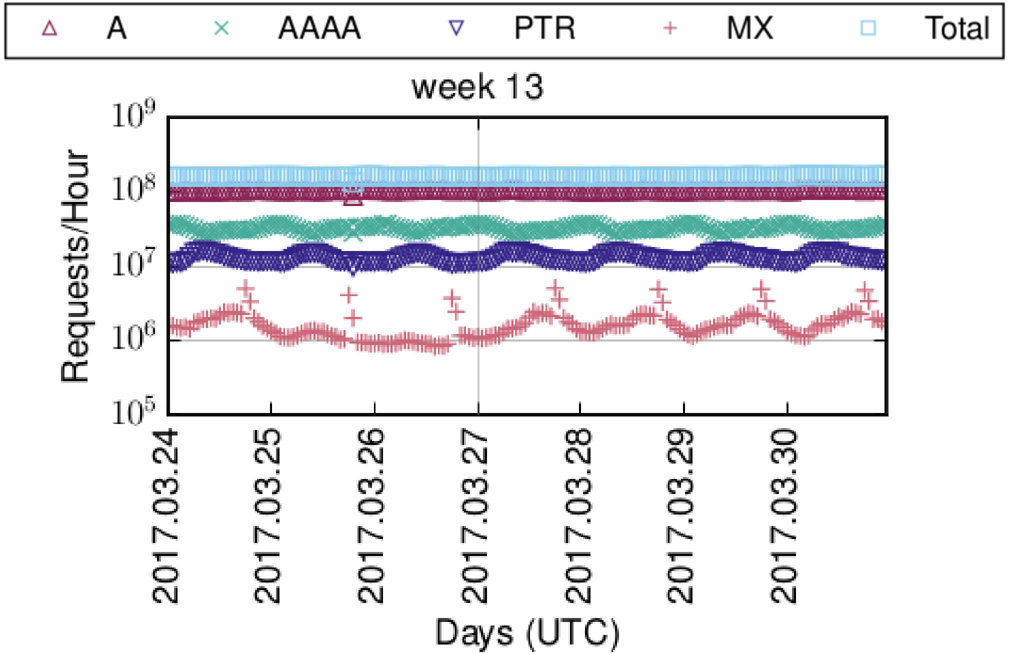

(a) Full Farsight dataset.

- Flat-line for A/Total
- Daily anti-pattern AAAA vs. PTR
  - ➢ Looks funny

Isolated biasing operator

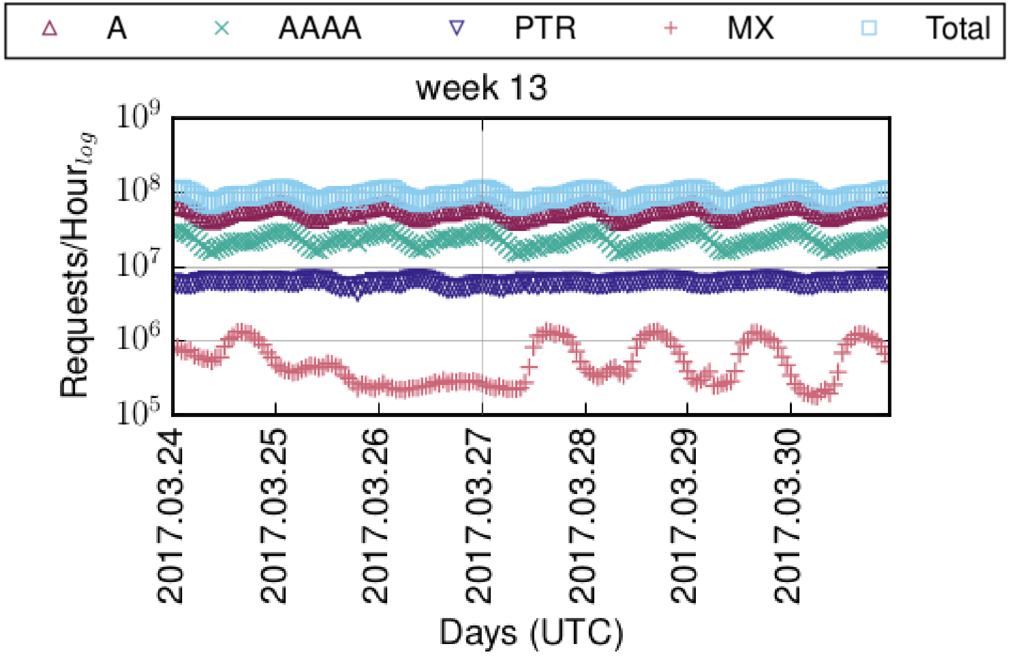

(b) Only biased operator.

- Cause: Single Operator with odd lookup pattern:
  - Flatline PTR
    - ip6.int. for 70::/8
    - DNS-SD for dell.com, apple.com etcetc. (~same No. Req/name)
- Close to 50% of the Dataset

TUDelft

Cleaned dataset

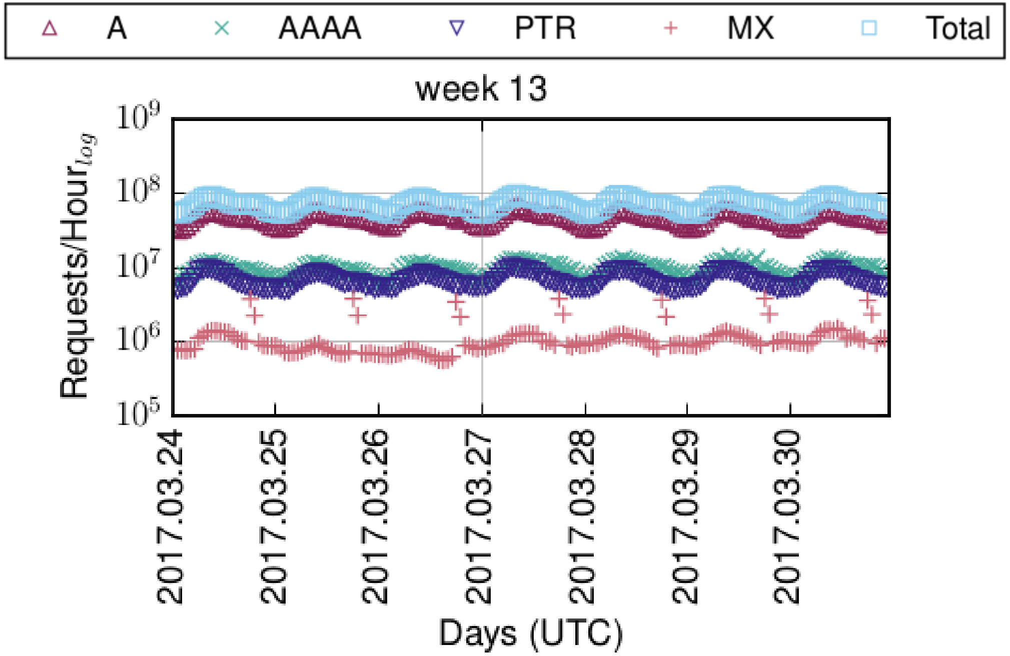

(c) w/o biased operator.

- Sollution: Filtering
- Patterns start to look as expected
- Outliers in MX: Single Russian ISP running a regular “Digest Mailinglist” for users

Types of PTR: v4, v6, DNS-SD

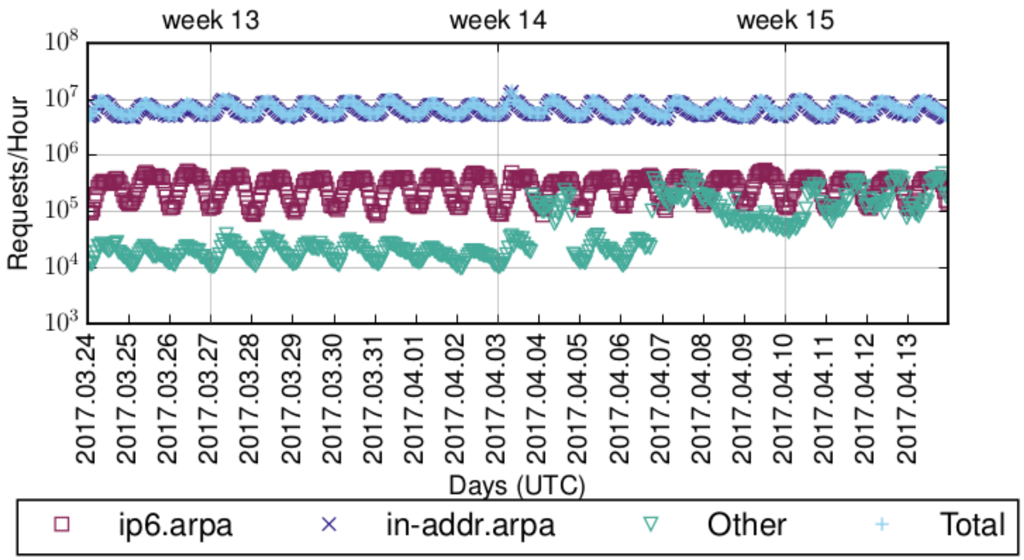

- Not all PTR are .arpa!
  - ~99% in-addr.arpa.
  - ~0.9% ip6.arpa
  - ~0.1% DNS-SD
- Outliers starting week 14:
  - Possible deployment of new set-top-box CPE (software)
  - Queries for TV Channel Domain

TUDelft

rDNS Response Codes: in-addr.arpa.

- Stable SERVFAIL socket (~3%)
- Relatively few NXDOMAIN (~25%)
- ~47% NOERROR
- ~15% REFUSED

rcode	in-addr.-arpa	ip6.arpa	ip6.arpa w/o Resv.
NOERROR	47.21%	4.00%	32.30%
NXDOMAIN	25.36%	94.87%	63.87%
REFUSED	15.47%	0.14%	1.11%
FAILURE	8.77%	0.81%	1.34%
SERVFAIL	3.17%	0.18%	1.38%
FORMERR	0.01%	≤0.01%	≤0.01%
NOTAUTH	≤0.01%	-	-
NOTIMP	≤0.01%	-	-

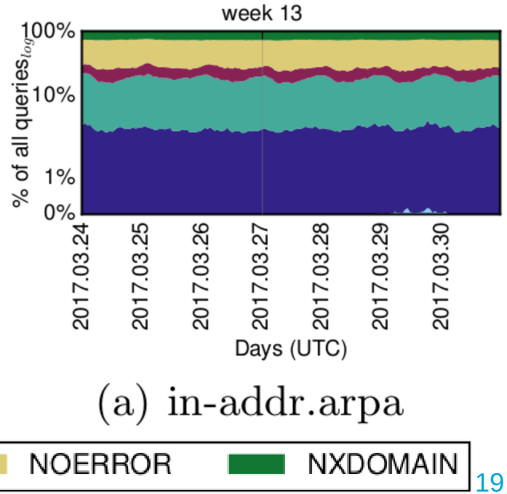

(a) in-addr.arpa

TUDelft

rDNS Response Codes: ip6.arpa.

- Hardly any SERVFAIL (>0.2%)
- Dominated by NXDOMAIN (~95%)
- Only 4% NOERROR
- Hardly any REFUSED (>0.2%)

rcode	in-addr.-arpa	ip6.arpa	ip6.arpa w/o Resv.
NOERROR	47.21%	4.00%	32.30%
NXDOMAIN	25.36%	94.87%	63.87%
REFUSED	15.47%	0.14%	1.11%
FAILURE	8.77%	0.81%	1.34%
SERVFAIL	3.17%	0.18%	1.38%
FORMERR	0.01%	≤0.01%	≤0.01%
NOTAUTH	≤0.01%	-	-
NOTIMP	≤0.01%	-	-

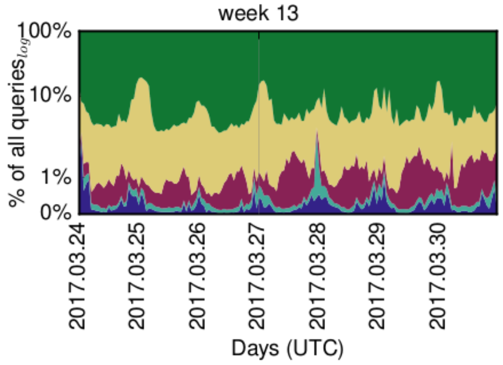

(b) ip6.arpa

rDNS Response Codes: ip6 w/o Resv.

- More SERVFAIL (>1.4%)
- Still strong on NXDOMAIN (~64%)
- Now ~32% NOERROR
- Still hardly any REFUSED (>1.2%)

rcode	in-addr.-arpa	ip6.arpa	ip6.arpa w/o Resv.
NOERROR	47.21%	4.00%	32.30%
NXDOMAIN	25.36%	94.87%	63.87%
REFUSED	15.47%	0.14%	1.11%
FAILURE	8.77%	0.81%	1.34%
SERVFAIL	3.17%	0.18%	1.38%
FORMERR	0.01%	$\leq$0.01%	$\leq$0.01%
NOTAUTH	$\leq$0.01%	-	-
NOTIMP	$\leq$0.01%	-	-

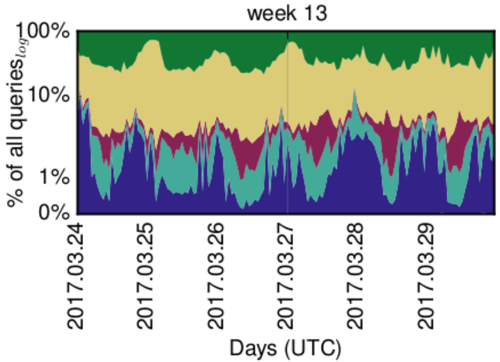

(c) ip6.arpa w/o Resv.

rDNS Response Types

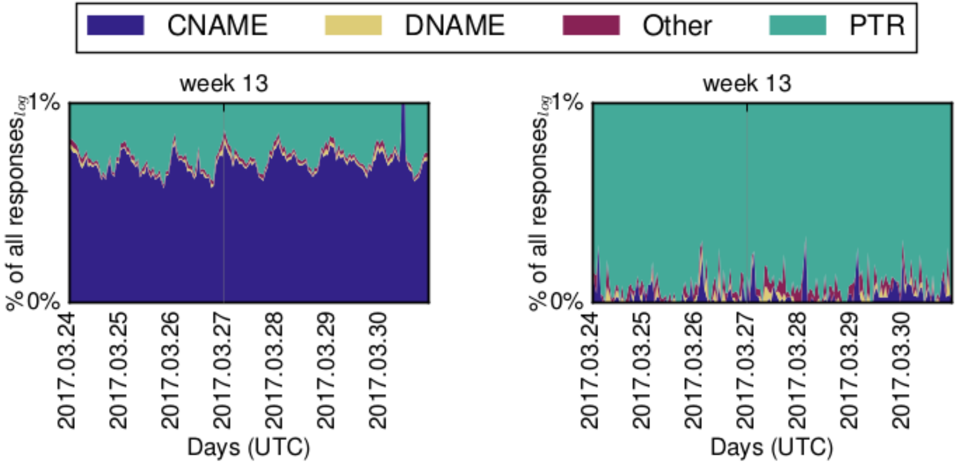

(a) in-addr.arpa

(b) ip6.arpa

- CNAMEs common for in-addr.arpa delegation
  - Hardly any in ip6.arpa.
- DNAMEs are a thing!

TU Delft

Passive Measurements Summary

- Beware of biases in data sets
- There is more v4 than v6 rDNS (100:1) and PTR ≠ rDNS
- Way more noise (priv./resv. For IPv6 rDNS)
- Less CNAMEs in v6 (as expected)
- More SERVFAIL in v4
- Less REFUSED in ip6.arpa. (Consistent with findings on lower IPv6 security, e.g., Czyz, Jakub, et al. “Don’t Forget to Lock the Back Door! A Characterization of IPv6 Network Security Policy.” *NDSS*. 2016.)

TU Delft

Active measurement…

TU Delft

Active rDNS measurements

- Easy for in-addr.arpa (brute-force)
- Hard for ip6.arpa (too large)
- Use RF8020 compliance as suggested in RFC7707 globally

Fiebig, Tobias, et al. "In rDNS We Trust: Revisiting a Common Data-Source's Reliability." *International Conference on Passive and Active Network Measurement*. Springer, Cham, 2018.

Enumerating (r)DNS trees

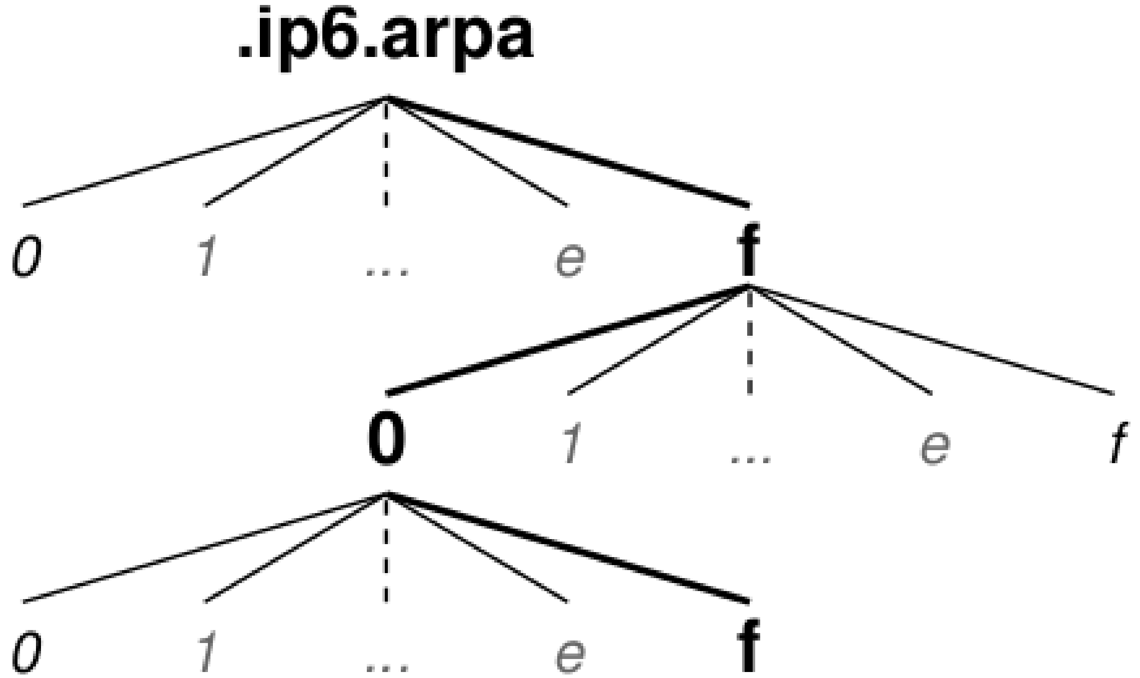

TUDelft

Collecting Data

- Used RFC8020 enumeration for v4 and v6
  - Quicker than brute-force for in-addr.arpa
    - Compared with a brute-force dataset

- Cluster of 16 machines (beware of the single IP stack)
  - Performed better than single machine for PAM2017 paper

TUDelft

Limitations

- We can not enumerate zone on RFC8020 violating authoritatives
- Cross-test with active trace:
  - 39.58% RFC8020 compliant
  - 46.42% always NXDOMAIN
  - 11.61% always return NOERROR
- Seeding makes things better, but we at best see only ~40%

TU Delft

rcodes in Active Measurements

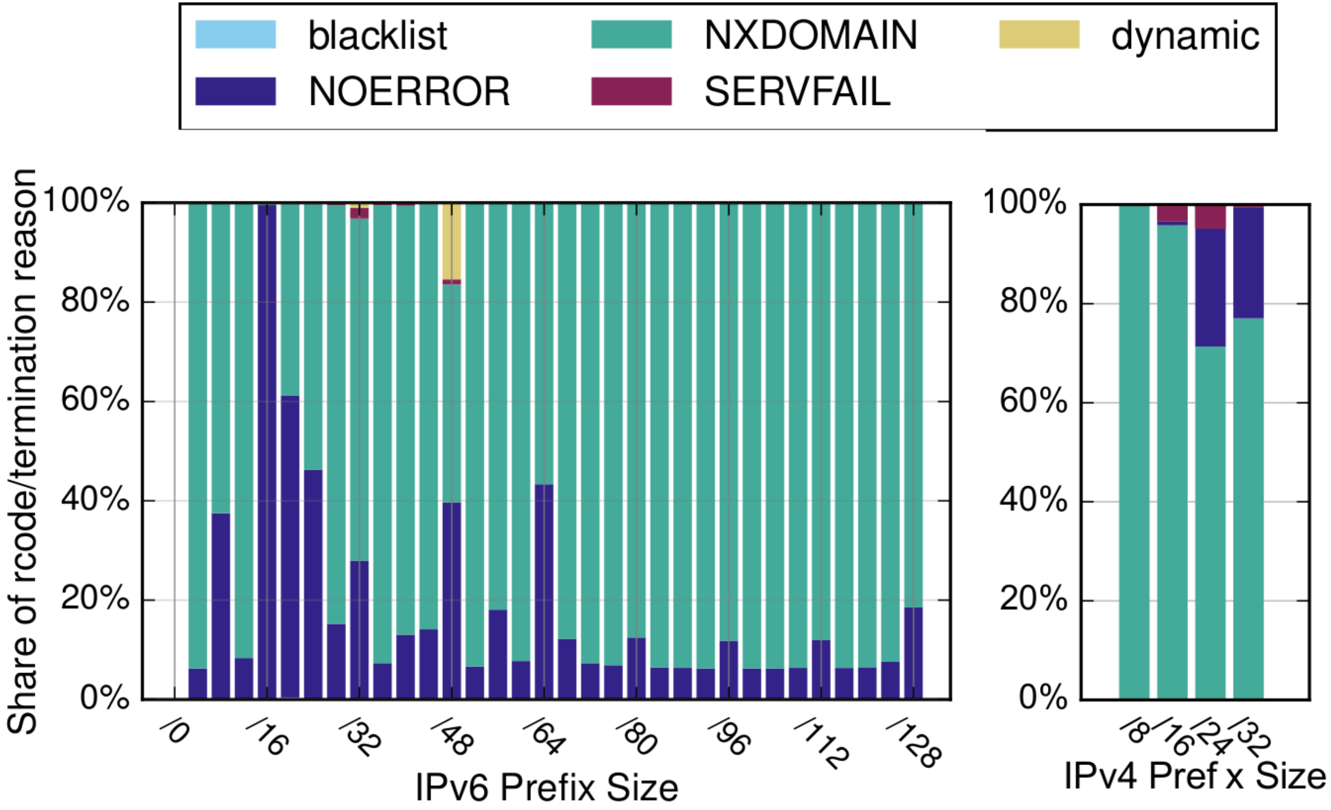

TUDelft

Delegations in rDNS

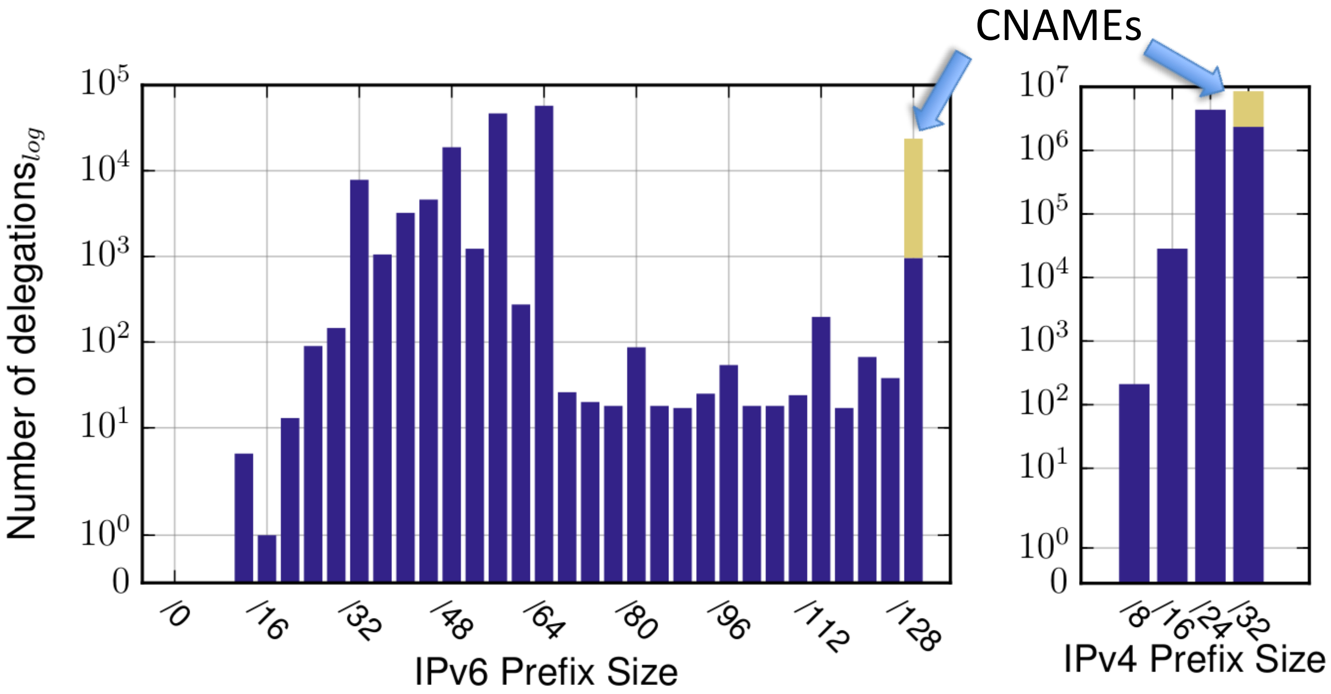

CNAMEs

- IPv4:
  - Mostly delegation for </24 zones (RFC2317)
- IPv6:
  - Heavy hitter: 87.81% of CNAMEs belong to a DHCPv6 setup (Dynamic Zone?)
  - 80.77% of the rest point to in-addr.arpa names
    - IPv4 first, consistent naming for multi-homed hosts

Special Case: rDNS64?

- Found a single operator mapping in-addr.arpa. to a /96 via CNAMEs
  - NAT64 range?

- Smart idea:
  - Preserves rDNS for customers
  - Does not break DNSSEC(!)
  - Should we have an RFC for this?

TUDelft

A/AAAA-less PTRs

- Found large operators with only PTR records set
  - Actual forward zones not populated or delegated?
  - Forward zones in split-view?
  - Potential information leak

TU Delft

Active Measurements Summary

- CNAME have different usecases in IPv4 and IPv6
- SERVFAIL is more common in v4 rDNS but overall relatively low
- IPv6 rDNS is top-driven
- Dynamically generated v6 zones are mostly /48
- We found a funny case of v4 rDNS for DNS64 delegation in a Japanese ISP
- There are names without a matching forward-record in .arpa

TU Delft

Summary & Conclusion

- PTR is not only for .arpa
- People still use rDNS
- IPv6 rDNS is ~1% of IPv4 rDNS
- We should take a look at whether people actually maintain their rDNS

TU Delft

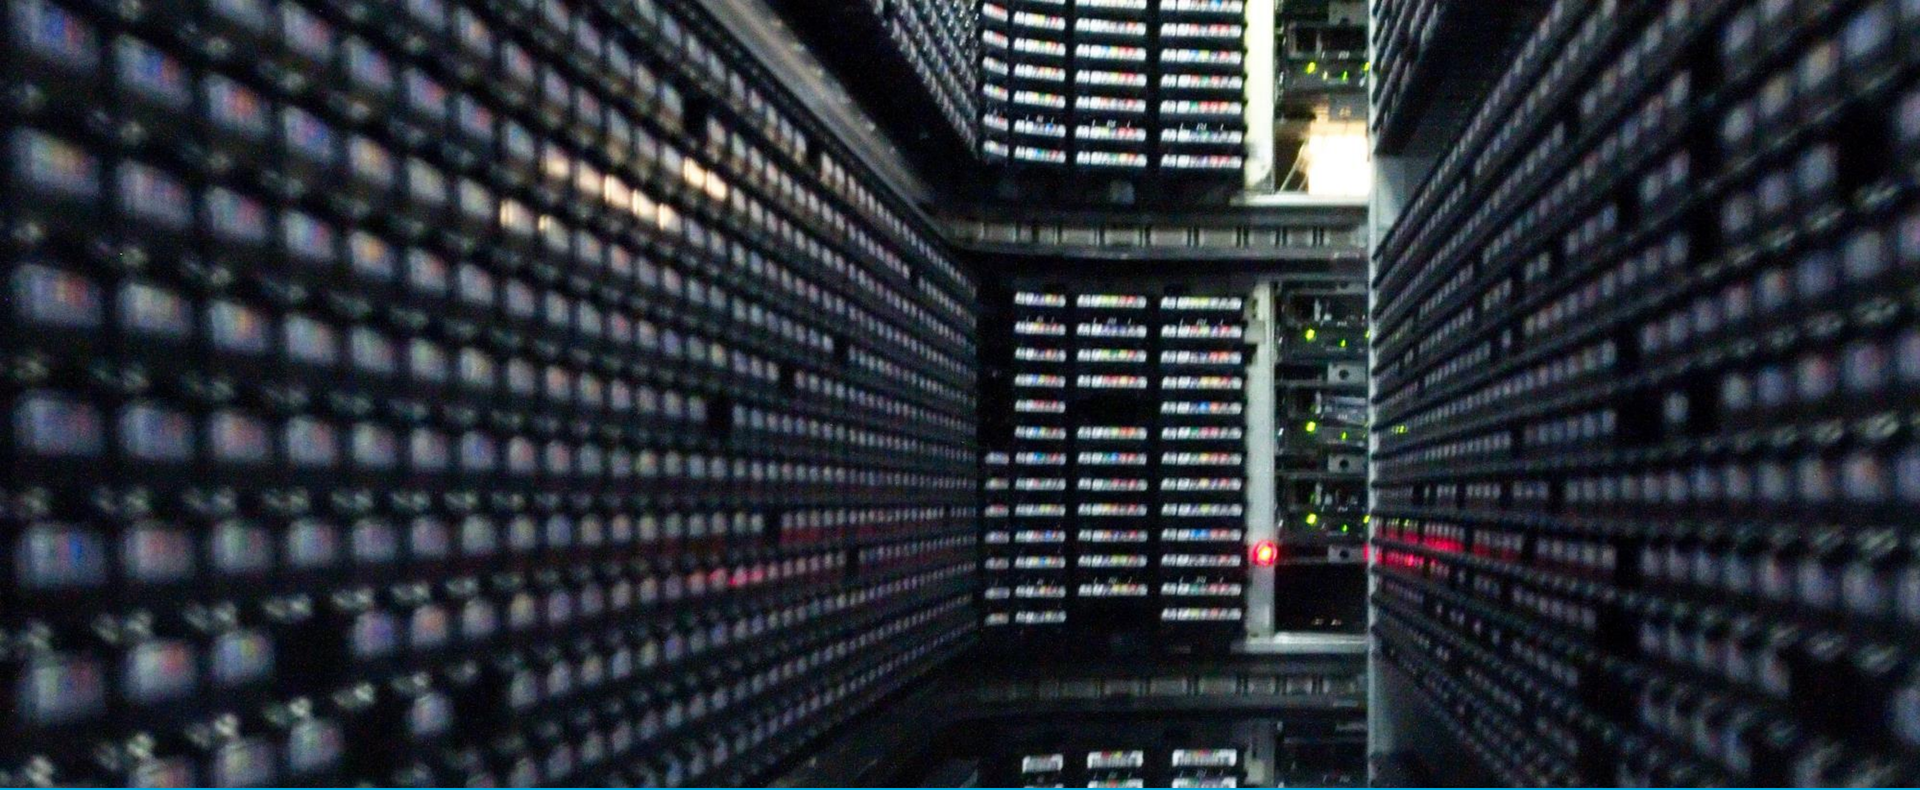

Backup Slides

TU Delft

rDNS Response Codes: Table

rcode	in-addr.-arpa	ip6.arpa	ip6.arpa w/o Resv.
NOERROR	47.21%	4.00%	32.30%
NXDOMAIN	25.36%	94.87%	63.87%
REFUSED	15.47%	0.14%	1.11%
FAILURE	8.77%	0.81%	1.34%
SERVFAIL	3.17%	0.18%	1.38%
FORMERR	0.01%	≤0.01%	≤0.01%
NOTAUTH	≤0.01%	-	-
NOTIMP	≤0.01%	-	-

TUDelft

Churn in Queried Names

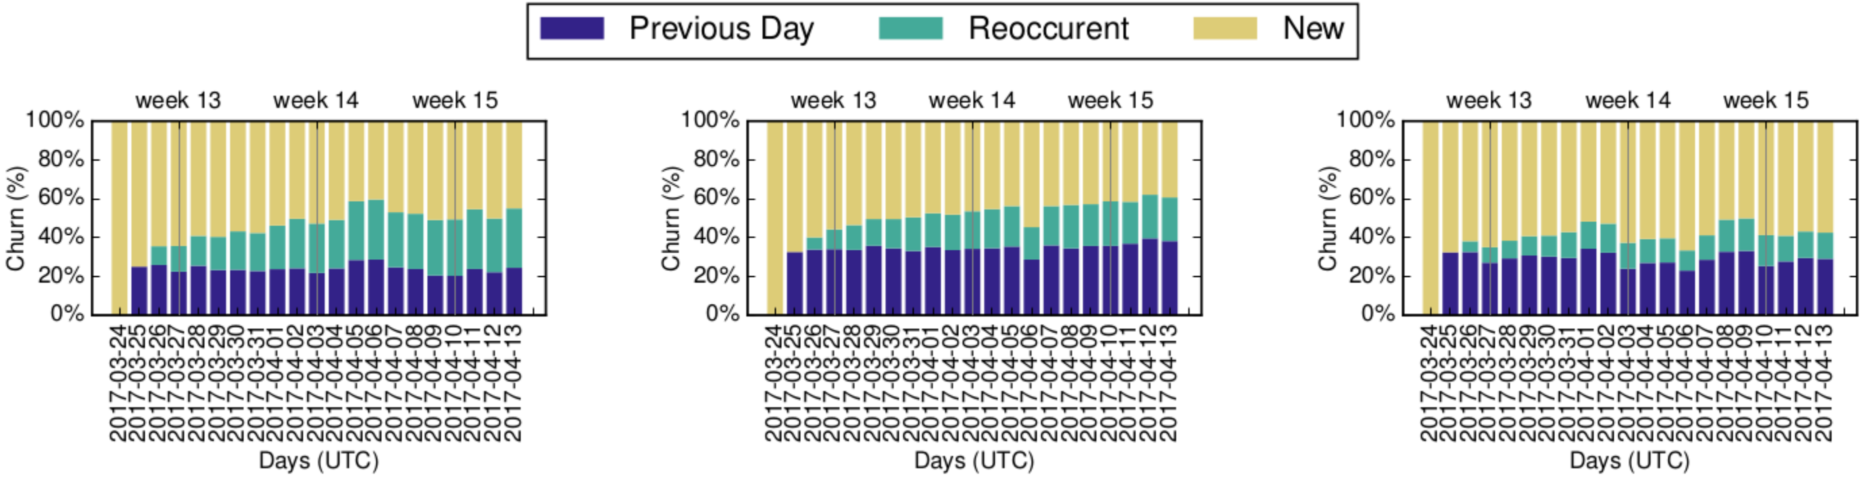

(a) in-addr.arpa (b) ip6.arpa (/64) (c) ip6.arpa (/128)

Biases and Volume

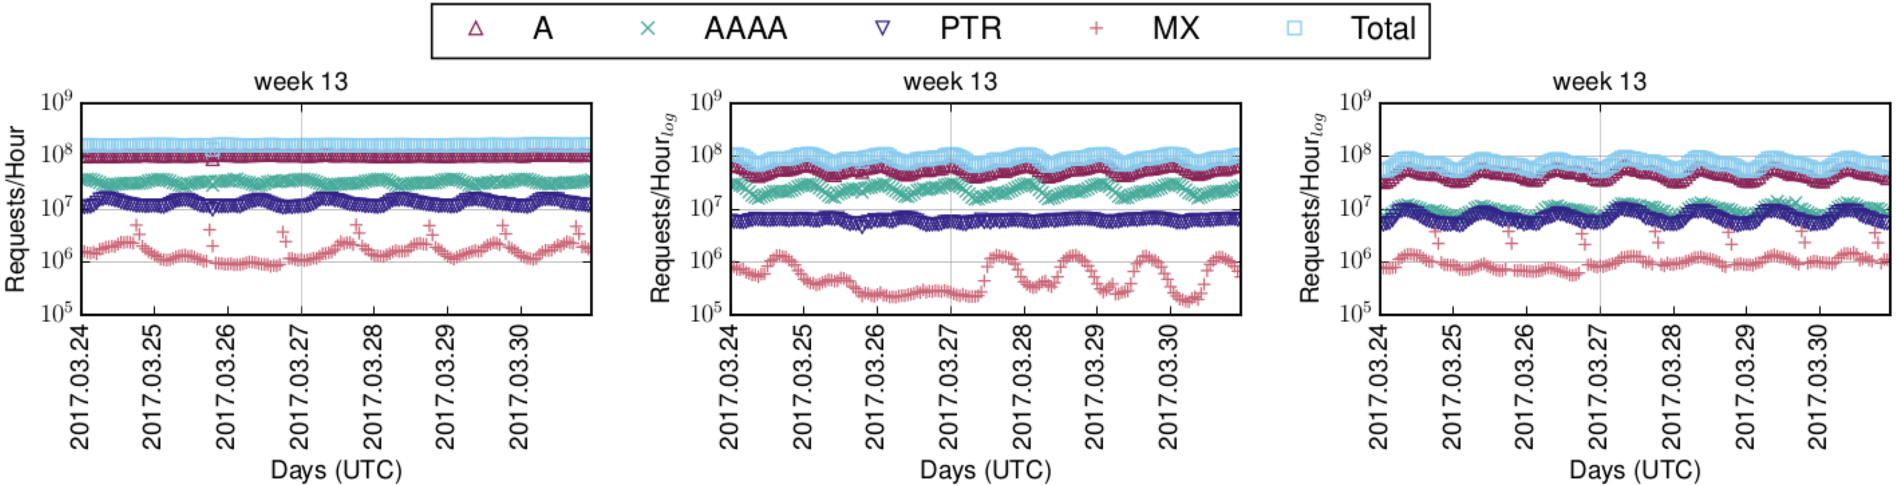

(a) Full Farsight dataset. (b) Only biased operator. (c) w/o biased operator.