



get the paper here



# How Ready is DNS for an IPv6-only World?

Florian Streibelt<sup>1</sup>, Patrick Sattler<sup>2</sup>, Franziska Lichtblau<sup>1</sup>,  
Carlos H. Gañán<sup>3</sup>, Anja Feldmann<sup>1</sup>, Oliver Gasser<sup>1</sup>, Tobias Fiebig<sup>1</sup>

<sup>1</sup> Max Planck Institute for Informatics, <sup>2</sup> TU München, <sup>3</sup> TU Delft

PAM 2023



max planck institut  
informatik

Technische  
Universität  
München



Paper:

<https://hdl.handle.net/21.11116/0000-000C-8817-1>

[https://link.springer.com/chapter/10.1007/978-3-031-28486-1\\_22](https://link.springer.com/chapter/10.1007/978-3-031-28486-1_22)

Code and Dataset:

<https://github.com/mutax/dns-v6-readiness>



# The IPv6-only Experience




**GEEKFLARE** TOOLS    Toolbox    Compiler    Log in    **Sign Up FREE**    Products ▾

**IPv6 Test**  
**en.wikipedia.org**

IP Address: 208.80.154.224    Test Time: Fri, Mar 17, 2023 3:34 PM (GMT +01:00)

Share Report  
Twitter Facebook LinkedIn WhatsApp

**Results**



Great, your site is accessible over IPv6.

**IPv6 address**

2620:0:862:ed1a::1

```
fls@glueball:~$ dig +short AAAA en.wikipedia.org
dyna.wikimedia.org.
2620:0:862:ed1a::1
fls@glueball:~$ ping6 -n -c 1 en.wikipedia.org
PING en.wikipedia.org(2620:0:862:ed1a::1) 56 data bytes
64 bytes from 2620:0:862:ed1a::1: icmp_seq=1 ttl=60 time=19.6 ms

--- en.wikipedia.org ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 19.629/19.629/19.629/0.000 ms
fls@glueball:~$
```



# The IPv6-only Experience



Hmm. We're having trouble finding that site.

We can't connect to the server at en.wikipedia.org.

If you entered the right address, you can:

- Try again later
- Check your network connection
- Check that Firefox has permission to access the web (you might be connected but behind a firewall)

Try Again



# Where does it break?

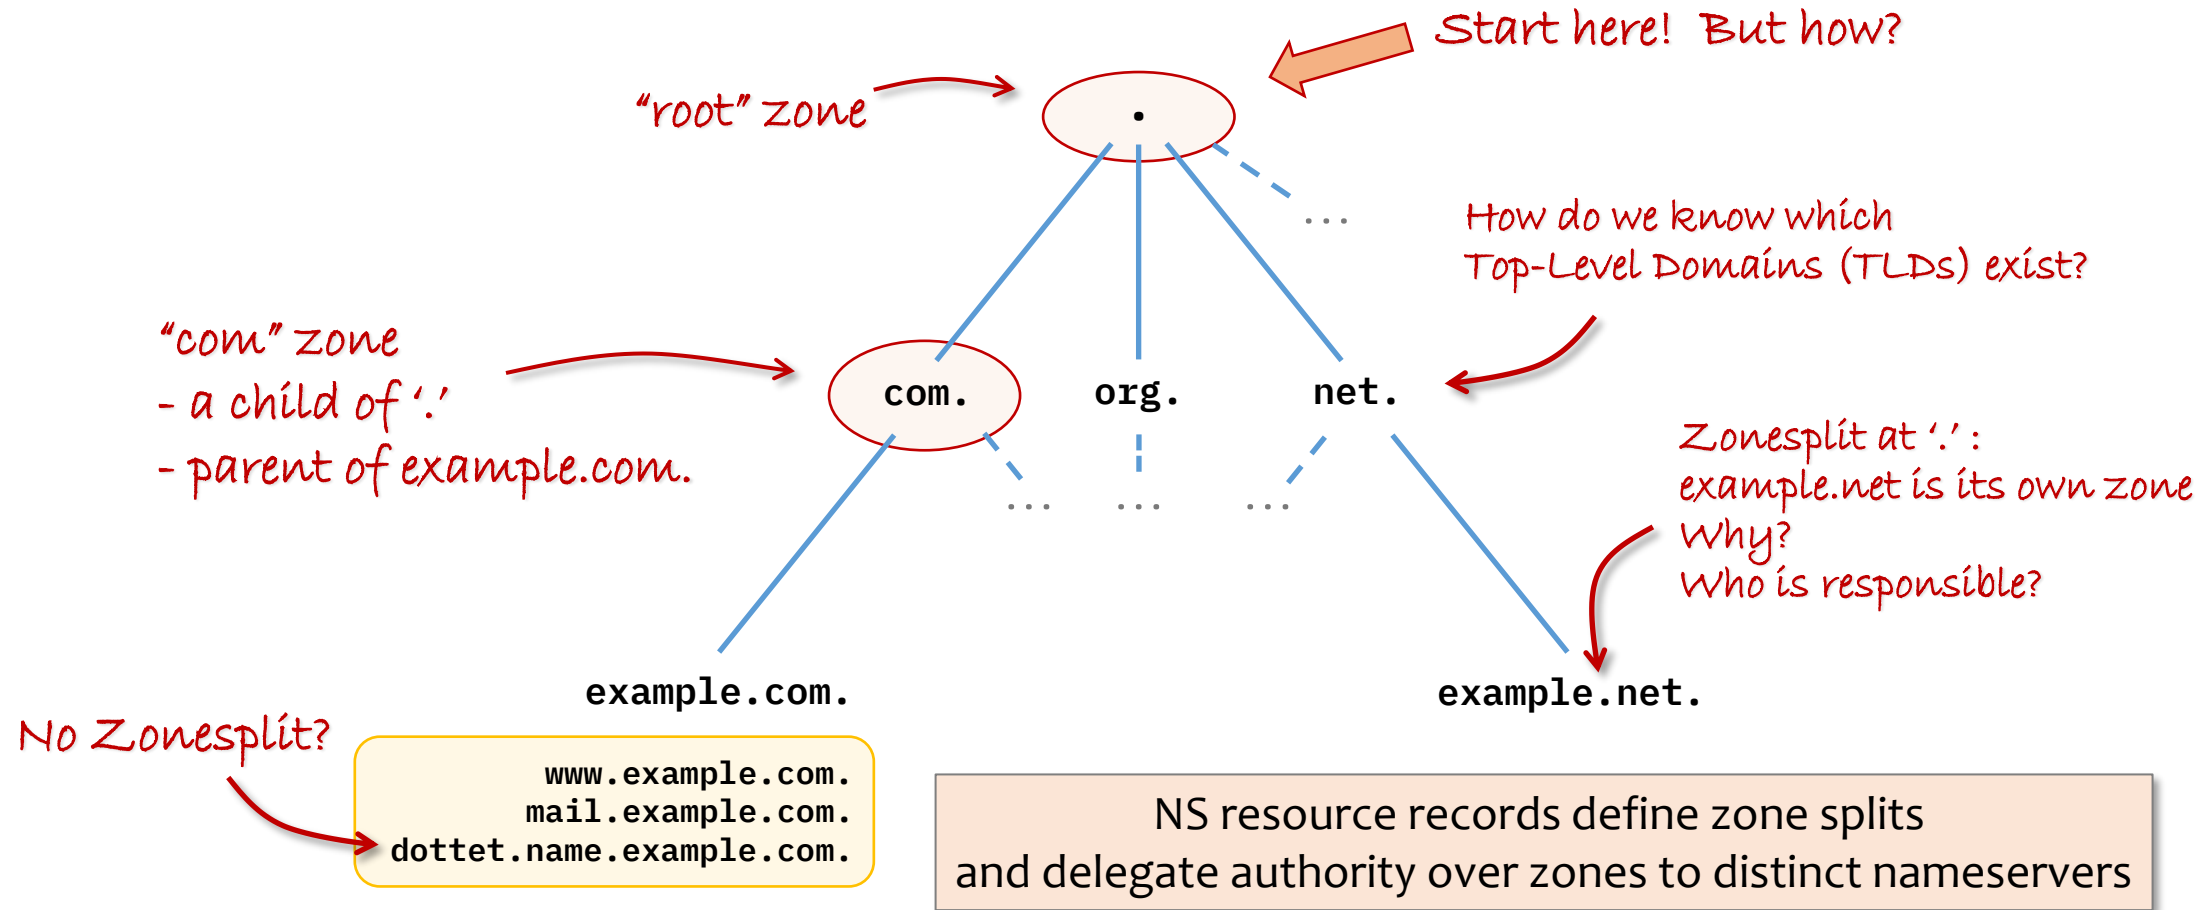


```
fls@glueball:~$ for NS in $(dig +short NS wikipedia.org); do  
> [ -z "$(dig +short AAAA ${NS})" ] && echo "No AAAA for $NS"  
> done  
No AAAA for ns1.wikimedia.org.  
No AAAA for ns2.wikimedia.org.  
No AAAA for ns0.wikimedia.org.  
fls@glueball:~$ █
```

We measure the current state of  
IPv6 resolvability in an IPv6-only scenario



# Let's talk about DNS...

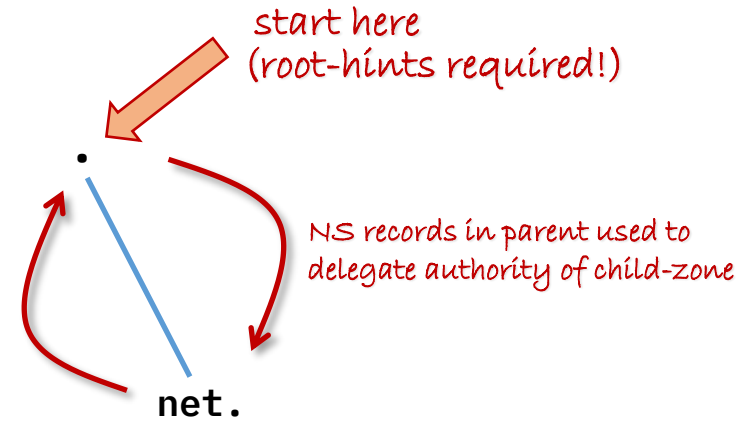


# How is it supposed to work?



root-hints (special case)	.	IN NS	a.root-servers.net.
	.	IN NS	b.root-servers.net.
	a.root-servers.net.	IN A	198.41.0.4
	a.root-servers.net.	IN AAAA	2001:503:ba3e::2:30
	b.root-servers.net.	IN A	199.9.14.201
	b.root-servers.net.	IN AAAA	2001:500:200::b
delegation	net.	IN NS	a.gtld-servers.net.
	net.	IN NS	b.gtld-servers.net.
GLUE-records	a.gtld-servers.net.	IN A	192.5.6.30
	a.gtld-servers.net.	IN AAAA	2001:503:a83e::2:30
	b.gtld-servers.net.	IN A	192.33.14.30
	b.gtld-servers.net.	IN AAAA	2001:503:231d::2:30

NS entries copied to parent  
GLUE records copied to parent  
out-of-band communication!



NS names that are "in-bailiwick"  
of their zones need GLUE in parent  
to break circular dependency

authoritative nameservers	net.	IN NS	a.gtld-servers.net.
	net.	IN NS	b.gtld-servers.net.

[ignoring TTLs and SOA (Start of Authority) on purpose]

Configuration of parent and child zone have to match,  
requires cooperation and coordination across organisations!





# What could possibly go wrong?

# No AAAA records for NS names



In short: “No nameserver has an IPv6 address”

*\$ORIGIN .com.*

*...*

*example.com. IN NS ns1.somedns.tld.*

*example.com. IN NS ns2.somedns.tld.*



*\$ORIGIN somedns.tld.*

*...*

*ns1 IN A 192.0.2.1*

*ns2 IN A 192.0.2.2*

*\$ORIGIN example.com.*

*@ IN NS ns1.somedns.tld.*

*@ IN NS ns2.somedns.tld.*

*www IN A 192.0.2.23*

*www IN AAAA 2001:db8::23*



# Missing GLUE in parent zone



```
$ORIGIN .com.
```

```
...
```

```
example.com.      IN NS  ns1.example.com.
```

```
example.com.      IN NS  ns2.example.com.
```

```
ns1.example.com.  IN A   192.0.2.1
```

```
ns2.example.com.  IN A   192.0.2.2
```

```
$ORIGIN example.com.
```

```
@      IN NS      ns1
```

```
@      IN NS      ns2
```

```
ns1    IN A       192.0.2.1
```

```
ns1    IN AAAA    2001:db8::1
```

```
ns2    IN A       192.0.2.2
```

```
ns2    IN AAAA    2001:db8::2
```

```
www    IN A       192.0.2.23
```

```
www    IN AAAA    2001:db8::23
```

```
$ dig +norecurse NS streibelt.net @i.gtld-servers.net.
```

```
; <<>> DiG 9.18.12 <<>> +norecurse NS streibelt.net @i.gtld-servers.net.
```

```
;; global options: +cmd
```

```
;; Got answer:
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 11227
```

```
;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 3, ADDITIONAL: 5
```

```
;; OPT PSEUDOSECTION:
```

```
; EDNS: version: 0, flags:; udp: 4096
```

```
;; QUESTION SECTION:
```

```
;streibelt.net.                IN      NS
```

```
;; AUTHORITY SECTION:
```

```
streibelt.net.      172800 IN      NS      ns01.streibelt.net.
```

```
streibelt.net.      172800 IN      NS      ns1.someserver.de.
```

```
streibelt.net.      172800 IN      NS      preon.streibelt.net.
```

```
;; ADDITIONAL SECTION:
```

```
ns01.streibelt.net. 172800 IN      A       164.68.125.157
```

```
ns01.streibelt.net. 172800 IN      AAAA    2a02:c207:3004:2982::1 ←
```

```
preon.streibelt.net. 172800 IN      A       188.68.54.227
```

```
preon.streibelt.net. 172800 IN      AAAA    2a03:4000:6:e08d::ffff ←
```

```
;; Query time: 3 msec
```

```
;; SERVER: 192.43.172.30#53(i.gtld-servers.net.) (UDP)
```

```
;; WHEN: Sun Mar 19 23:51:01 CET 2023
```

```
;; MSG SIZE rcvd: 200
```



# No AAAA for in-bailiwick NS names



In short: Mismatch between Zones

*\$ORIGIN .com.*

...

*example.com. IN NS ns1.example.com.*

*example.com. IN NS ns2.example.com.*

*ns1.example.com. IN A 192.0.2.1*

→ *ns1.example.com. IN AAAA 2001:db8::1*

*ns2.example.com. IN A 192.0.2.2*

→ *ns2.example.com. IN AAAA 2001:db8::2*

*\$ORIGIN example.com.*

*@ IN NS ns1*

*@ IN NS ns2*

*ns1 IN A 192.0.2.1*

*ns2 IN A 192.0.2.2*

*www IN A 192.0.2.23*

*www IN AAAA 2001:db8::23*

Triggers security feature implemented in some resolvers,  
e.g. Unbound with *harden-glue: yes* (the default)



# Zone of out-of-bailiwick NSes not resolving



In short: The zones of the nameservers “have to work”

```
$ORIGIN .com.  
...  
example.com.  IN NS  ns1.somedns.tld.  
example.com.  IN NS  ns2.somedns.tld.
```

```
$ORIGIN example.com.  
  
@      IN NS  ns1.somedns.tld.  
@      IN NS  ns2.somedns.tld.  
  
www    IN A    192.0.2.23  
www    IN AAAA 2001:db8::23
```

```
$ORIGIN somedns.tld.  
...  
@      IN NS  ns1.ipv4only.tld  
@      IN NS  ns2.ipv4only.tld
```

```
ns1    IN A    192.0.2.1  
ns1    IN AAAA 2001:db8::1  
ns2    IN A    192.0.2.2  
ns2    IN AAAA 2001:db8::2
```



# Parent zone not IPv6-resolvable



In short: all parent zones “have to work”

*\$ORIGIN example.com.*

@	IN NS	ns1
@	IN NS	ns2
ns1	IN A	192.0.2.1
ns1	IN AAAA	2001:db8::1
ns2	IN A	192.0.2.2
ns2	IN AAAA	2001:db8::2
sub	IN NS	ns1.sub
sub	IN NS	ns2.sub
ns1.sub	IN A	192.0.2.21
ns1.sub	IN AAAA	2001:db8::21
ns2.sub	IN A	192.0.2.22
ns2.sub	IN AAAA	2001:db8::22

*\$ORIGIN .com.*

...

example.com.	IN NS	ns1.example.com.
example.com.	IN NS	ns2.example.com.
ns1.example.com.	IN A	192.0.2.1
ns2.example.com.	IN A	192.0.2.2

*\$ORIGIN sub.example.com.*

@	IN NS	ns1
@	IN NS	ns2
ns1	IN A	192.0.2.21
ns1	IN AAAA	2001:db8::21
ns2	IN A	192.0.2.22
ns2	IN AAAA	2001:db8::22
www	IN A	192.0.2.122
www	IN AAAA	2001:db8::122



# The five Horseman of DNS Misconfiguration



The presented issues are:

- not IPv6-specific,  
but usually quickly noticed in IPv4!
- not mutually exclusive



Original by William Bramhall, NYDN

One misconfigured zone  
will break all it's child zones.  
And the child zones. ...

# Datasets



## Passive Dataset: Farsight SIE

- Coverage: global
- Cache misses of recursors
- January 2015  
until August 2022

## Passive Dataset: Zonefiles

- .com, .net, and other gTLDs  
(starting mid of 2016)
- ICANN Centralized Zone Data Service  
for TLDs  
(from April 2017 onward)

Additionally for the coverage analysis:

- Zone file data from .se, .nu, and .ch  
that are publicly available

We compare Farsight's data  
to more than 1.1k zones as of August 2022

## Active Measurements

- Alexa Top 1 M,  
Aug 2022 (498k)
- One VP
- 56 M queries
- Oct 11-14 & 22-24
- Dataset publicly available



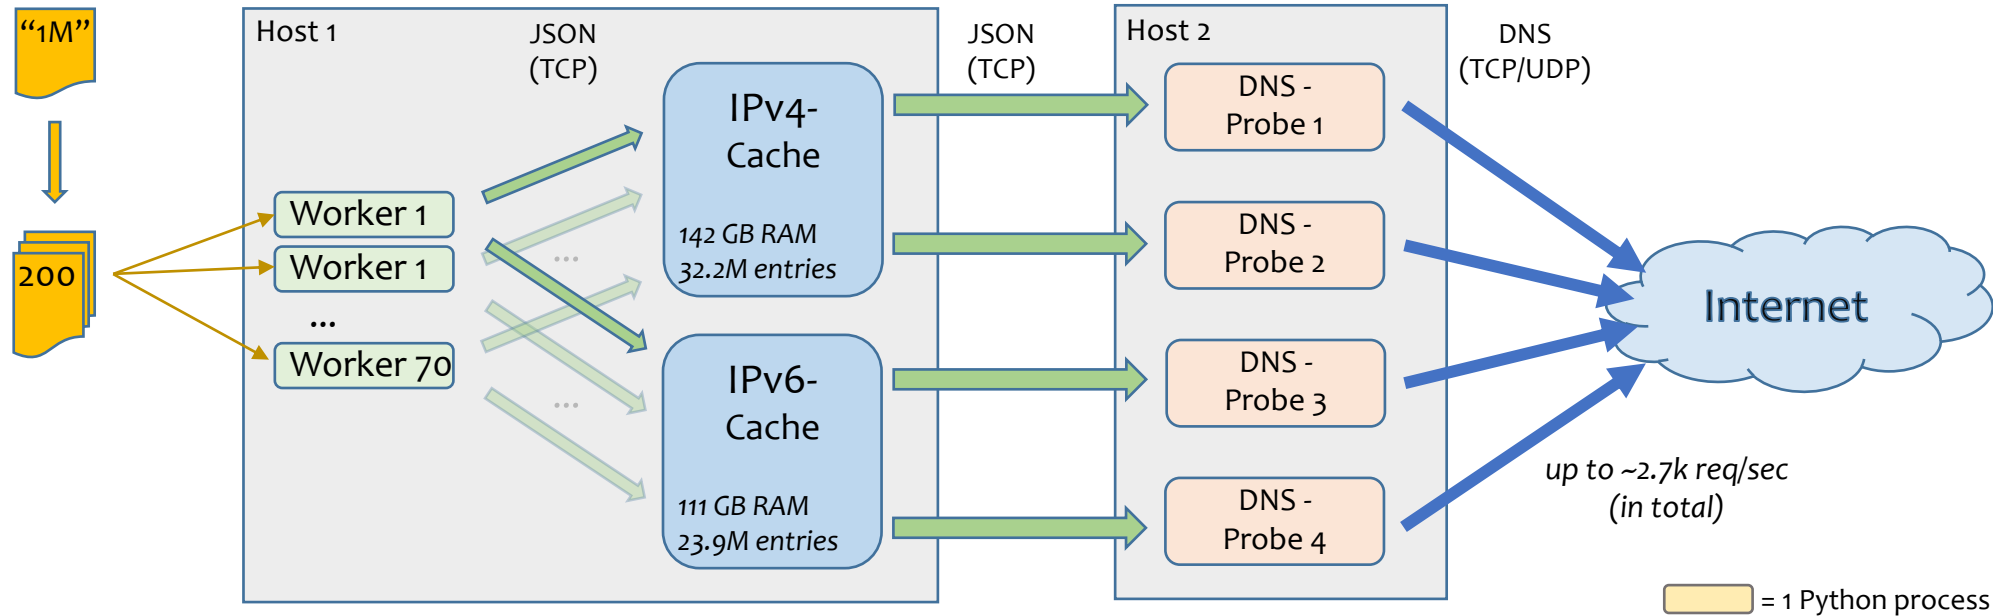
# Methodology



- Common for all:  
Do name resolution starting from the root
- Why can we do that?
  - Farsight dataset:
    - Clients ask for A and AAAA (happy eyeballs)
    - GLUE records contain A and AAAA
  - Zone files:
    - Ground truth, data should be complete
- Active Measurements
  - Query all authoritative nameservers, combine responses
  - Don't break stuff and be careful



# Active Measurements: System Design



Alexa list split in chunks of 200 entries,  
consumed by 70 workers in parallel,  
each worker starts fresh from root,

workers & caches co-located

Caches:  
mark failed servers,  
store error responses,  
content available as  
public dataset

Custom TCP protocol,  
based on plain JSON,  
Nagle disabled,  
persistent connection

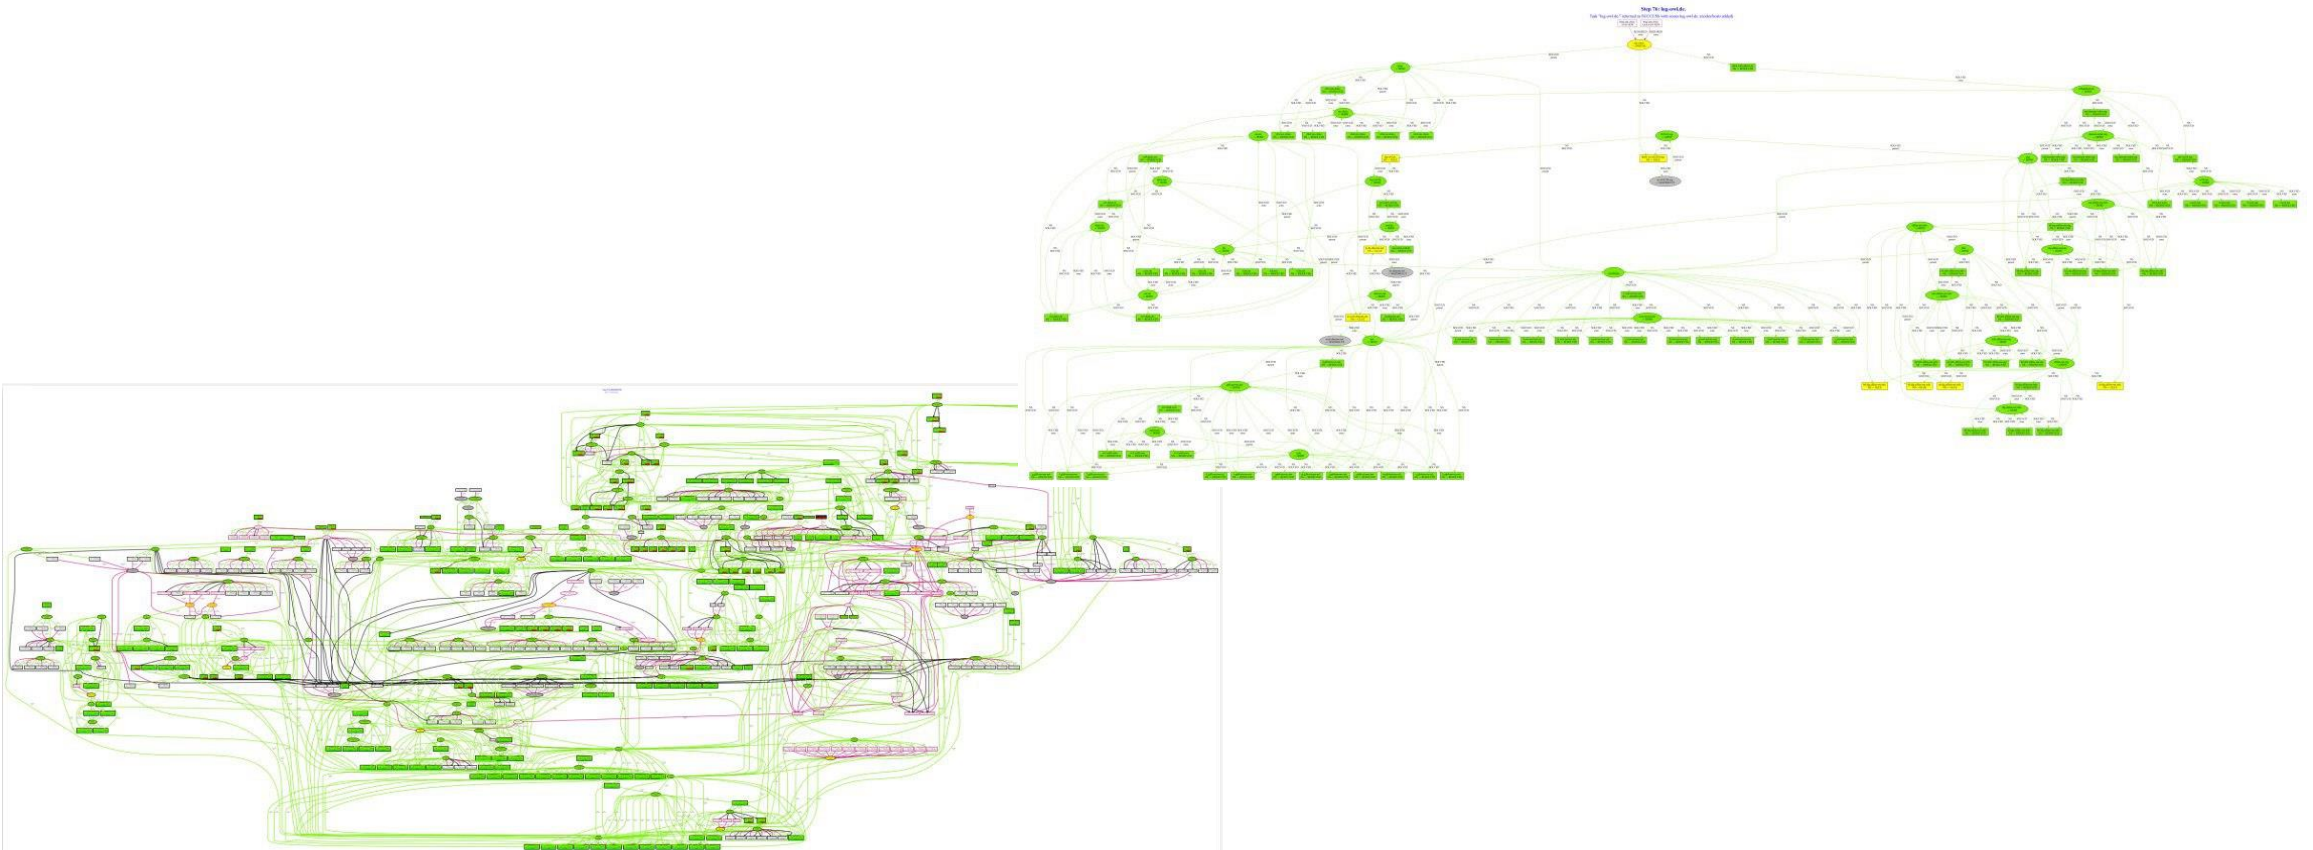
Probes:  
rate limiting,  
retries,  
truncation/TCP  
no-EDNS fallback



# Results and some Modern Art



Comprehensive visualisation is almost impossible...



# With Time comes... IPv6?

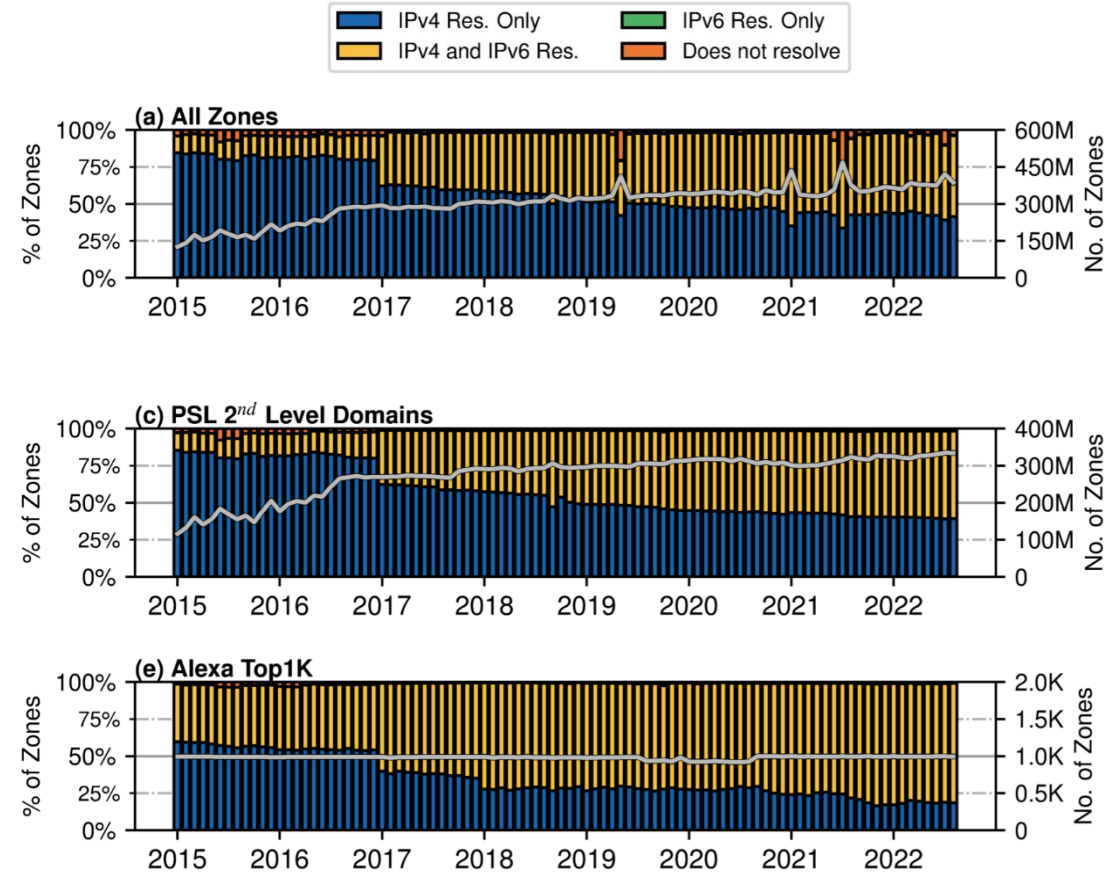


Fig. 3: Per month: # of zones (gray line-right y-axis) and IPv4/IPv6 resolvability in % (left y-axis).



# Why do Zones (SERV) Fail?

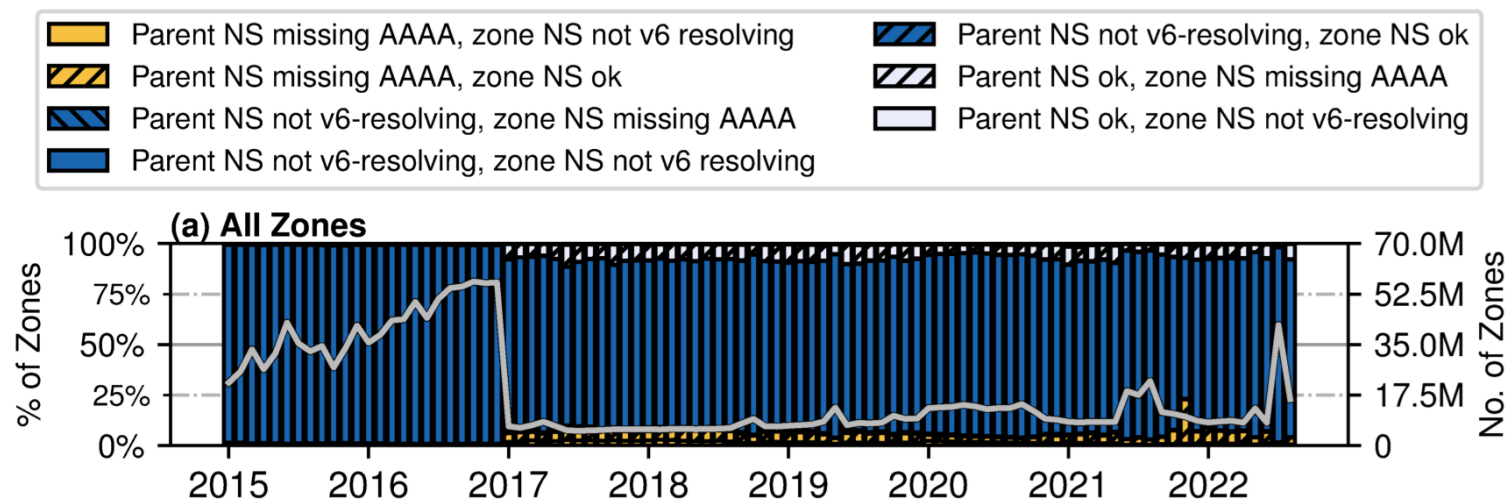
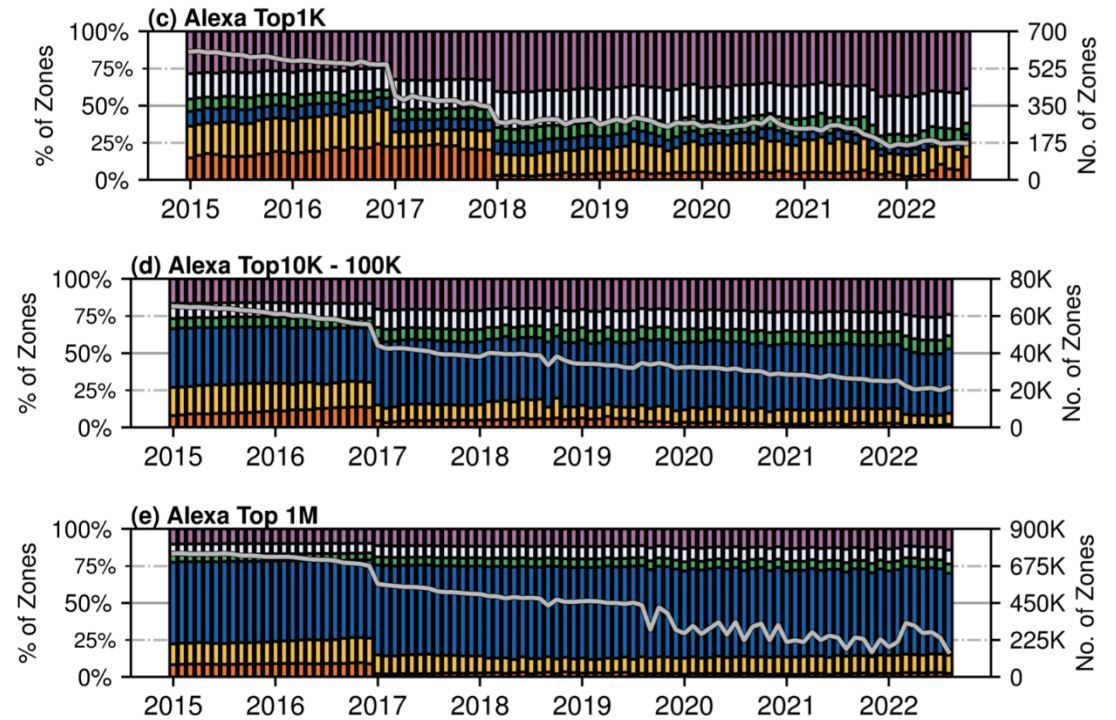
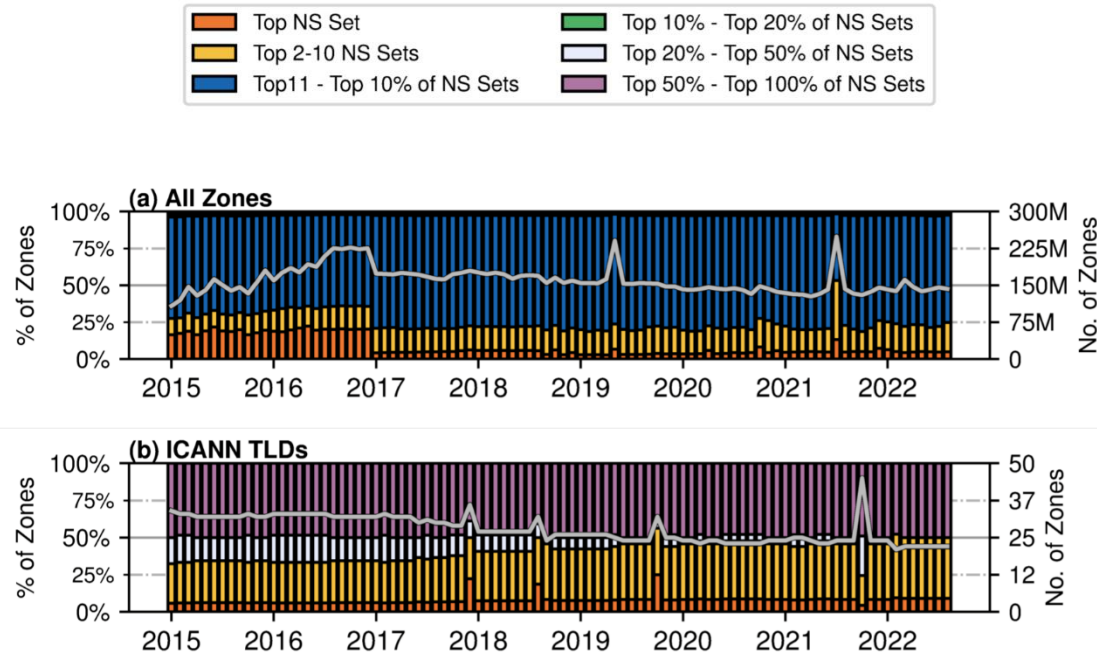


Fig. 5: Per month: # of zones not IPv6-resolvable with AAAA or GLUE for NS (gray line-right y-axis) and causes for IPv6 resolution failure in % (left y-axis).

The majority of zones is not IPv6-only resolvable  
because their parent zone(s) fail already



# Centralization: Opportunity and Risk?



A very small set of nameservers (and thus operators)  
runs most of the DNS infrastructure



# Misconfiguration hidden in plain sight!



- Misconfiguration, Hard- or Software failures are common
- Often unnoticed for a long time – why?
- DNS hides problems:
  - Resilience was a design goal
  - Good for user experience, bad for ops
- Monitoring needs to take that into account

Is this sufficient?

```
# dig www.google.com && echo "DNS works!"
```



# Some Caveats



- What we measure can have different reasons
  - Misconfiguration, e.g. oversight, inexperienced operators
  - Deliberate choice, e.g. during migration, maintenance
  - Bugs in our code, setup, vendor gear, ...
  - Rate limiting, e.g. no response vs. edns/truncation/...
- DNS is a complicated protocol!
  - Dependency loops!
  - Corner cases!
  - Many extensions!

Are we supposed to know the protocols better  
than the RFC authors and operators?  
Get help! Talk to people!



# Summary & Future Work



## Summary

- Passive measurement study with root cause analysis for broken IPv6 delegation in an **IPv6-only** setting
- Confirmation via active measurements
- August 2022: 44.9% of considered zones not IPv6-only resolvable  
Most common: zone or parent NS unresolvable
- Recommendation: Monitor IPv6 across entire delegation chain
- IPv6 readiness of the web may be impaired by non IPv6 resolvable domains

## Future work

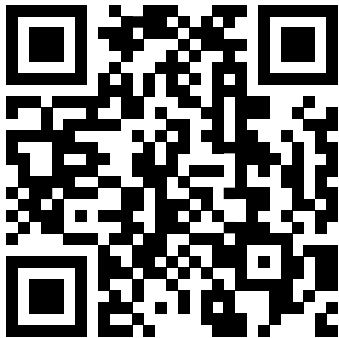
- Continuous scans of IPv6 resolvability of web assets to raise awareness
- Code polishing (this is a PoC) - implementation will be open sourced
- We will provide a reduced toolset for operators



# Free Stuff!

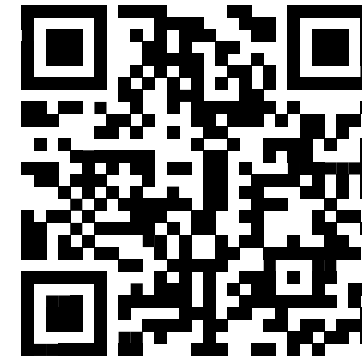


*Scan this to get our paper*



<https://hdl.handle.net/21.11116/0000-000C-8817-1>  
[https://link.springer.com/chapter/10.1007/978-3-031-28486-1\\_22](https://link.springer.com/chapter/10.1007/978-3-031-28486-1_22)

*Scan that to get datasets and code*



<https://github.com/mutax/dns-v6-readyness>



# Active Measurements, Requests/Day



Cache entries added per day and IP protocol

Date	IPv4 entries	IPv6 entries
2022-10-10	8355	7056
2022-10-11	559871	402949
2022-10-12	7246714	5265976
2022-10-13	12644623	9717726
2022-10-14	5258726	4097060
2022-10-22	5074976	3416827
2022-10-23	1283052	909097
2022-10-24	141556	102539

Dataset linked at <https://github.com/mutax/dns-v6-readiness>



# Farsight mode of operation

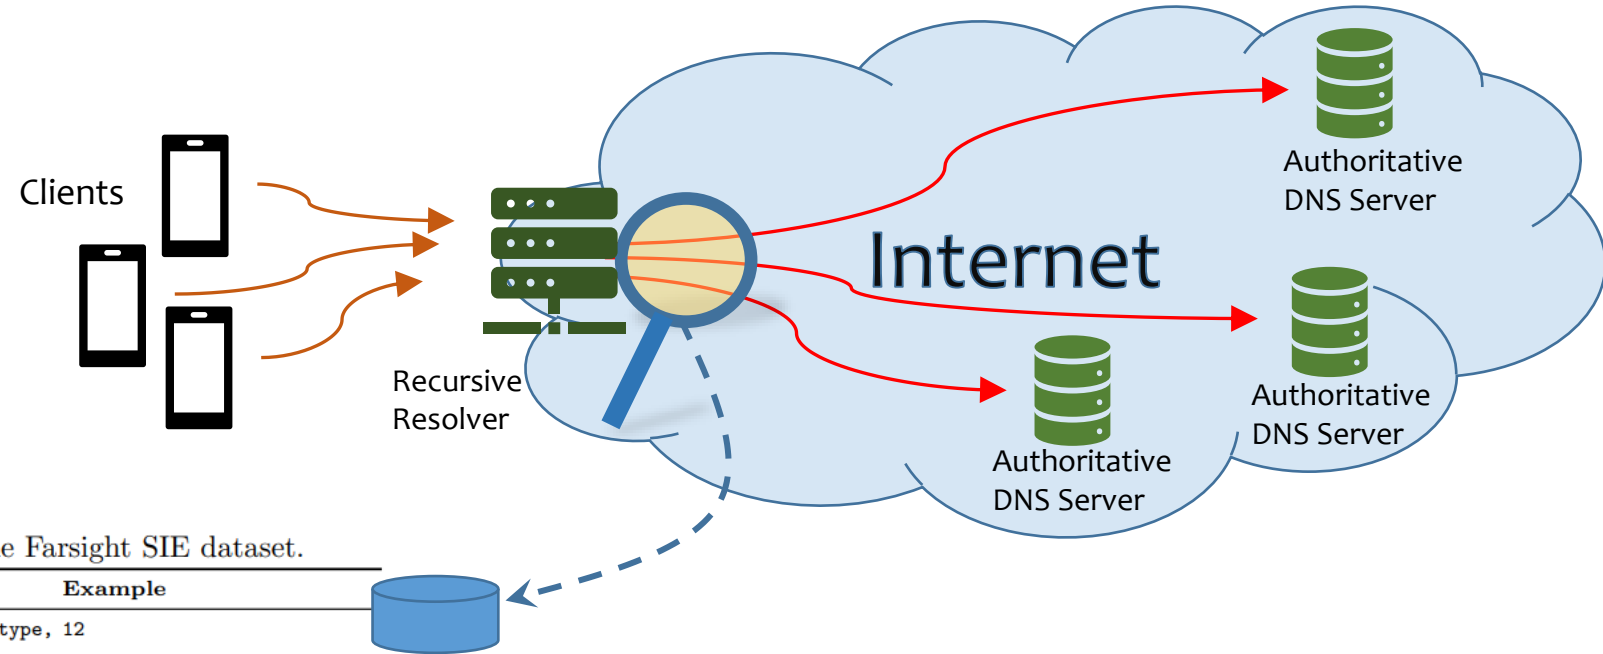


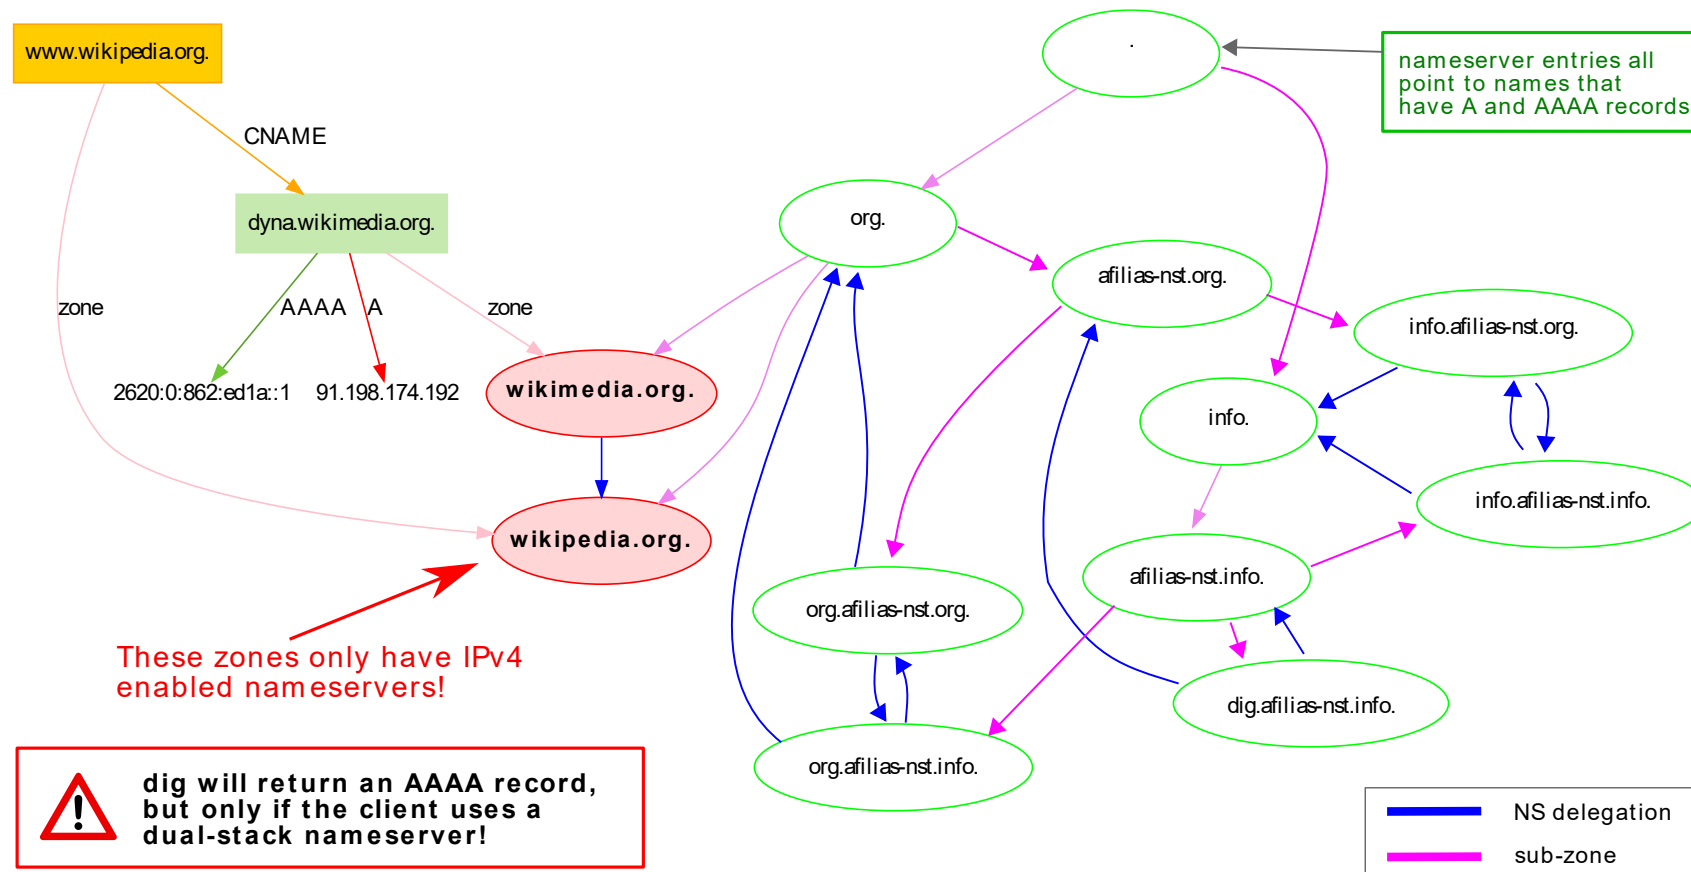
Table 1: List of data fields in the Farsight SIE dataset.

Field	Description	Example
count	# of times the tuple <rrname, rrtype, 12 bailiwick, rdata> has been seen.	
time_first	Unix timestamp of the first occurrence of the unique tuple during the data slice	1422251650
time_last	Unix timestamp of the last occurrence of the unique tuple during the data slice.	1422251650
rrname	Requested name in the DNS.	example.com
rrtype	Requested RRtype of the query.	NS
bailiwick	Zone authoritative for the reply.	com
rdata	List of all responses received in a single query.	["ns1.example.com", "ns2.example.com"]

No individual queries or client addresses logged, only cache-misses, analyzed as monthly aggregates



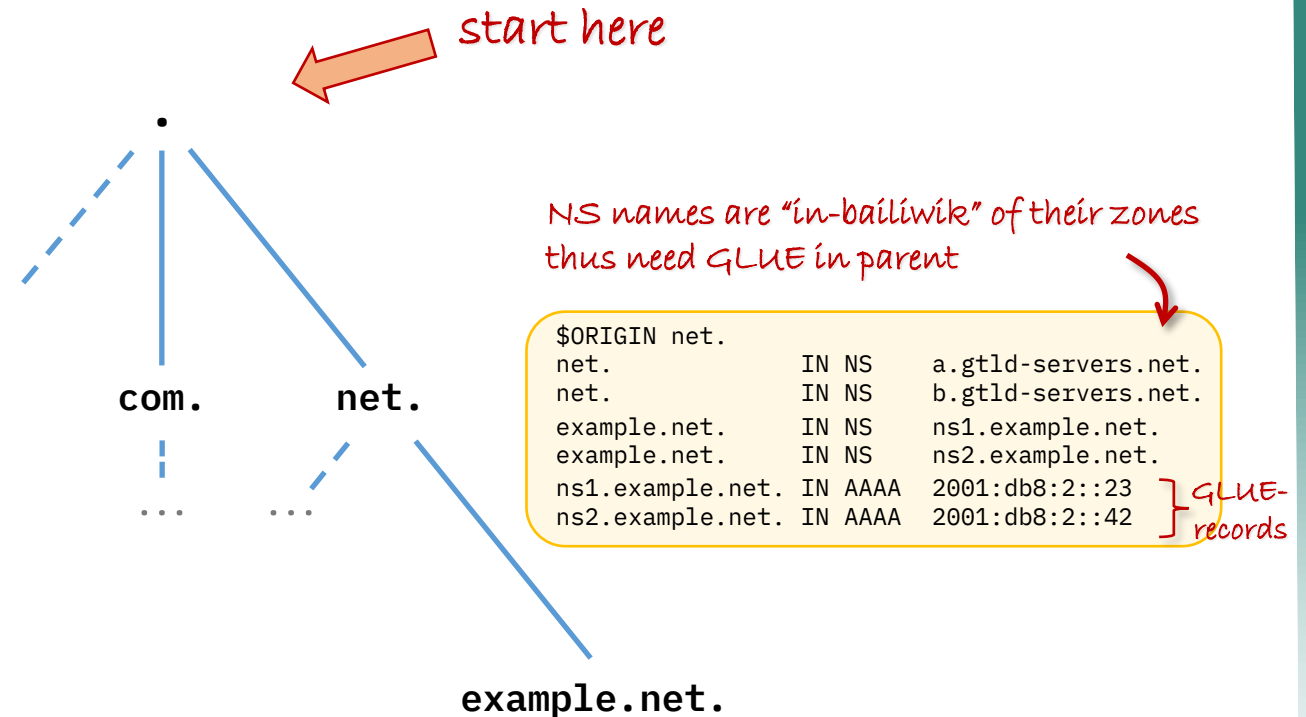
# Zooming in: www.wikipedia.org



# Name resolution, step-by-step



root-hints	\$ORIGIN .		
	@	IN NS	a.root-servers.net.
	@	IN NS	b.root-servers.net.
	a.root-servers.net.	IN A	198.41.0.4
	a.root-servers.net.	IN AAAA	2001:503:ba3e::2:30
delegation	b.root-servers.net.	IN A	199.9.14.201
	b.root-servers.net.	IN AAAA	2001:500:200::b
GLUE-records	net.	IN NS	a.gtld-servers.net.
	net.	IN NS	b.gtld-servers.net.
	a.gtld-servers.net.	IN A	192.5.6.30
	a.gtld-servers.net.	IN AAAA	2001:503:a83e::2:30
	b.gtld-servers.net.	IN A	192.33.14.30
	b.gtld-servers.net.	IN AAAA	2001:503:231d::2:30



```
$ORIGIN example.com.
@      IN NS    ns1.example.net.
@      IN NS    ns2.example.net.
...
```

"out-of-bailiwick",  
no GLUE needed, but additional queries

```
$ORIGIN example.net.
@      IN NS    ns1.example.net.
@      IN NS    ns2.example.net.
ns1    IN AAAA   2001:db8:2::23
ns2    IN AAAA   2001:db8:2::42
www    IN AAAA   2001:db8:2::80
```

authoritative nameservers

Task: resolve `www.example.net`



# Active Measurements and Passive Datasets align



- We find huge overlap in our datasets
- Todo: Active measurements confirm evaluation

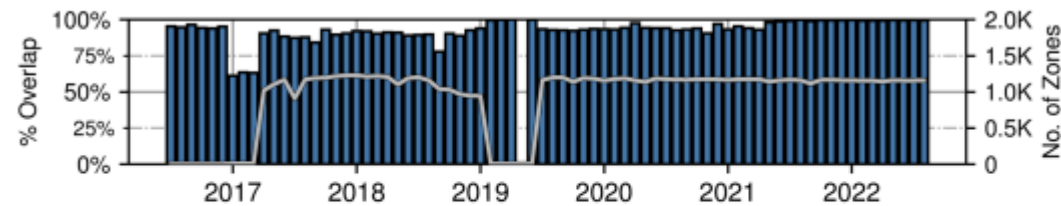


Fig. 2: Zone coverage of Farsight data and number of zones used for the evaluation. We used available zone files to determine the share of covered second level domains by Farsight's dataset. Please note the dip in the graph from February to August 2019, where our zone file collection was limited, i.e., we only collected few zones with high coverage (February - April and July, including .com), or no data at all (May and June).



# Datasets



## Passive Dataset: Farsight SIE

- Coverage: global
- Cache misses only
- Monthly aggregates
- January 2015  
until August 2022

## Passive Dataset: Zonefiles

- .com, .net,  
and other gTLD zone files  
(starting mid of 2016)
- ICANN Centralized Zone Data Service  
zone file data for all available TLDs  
(from April 2017 onward)

Additionally for the coverage analysis:

- Zone file data from .se, .nu, and .ch  
that are publicly available

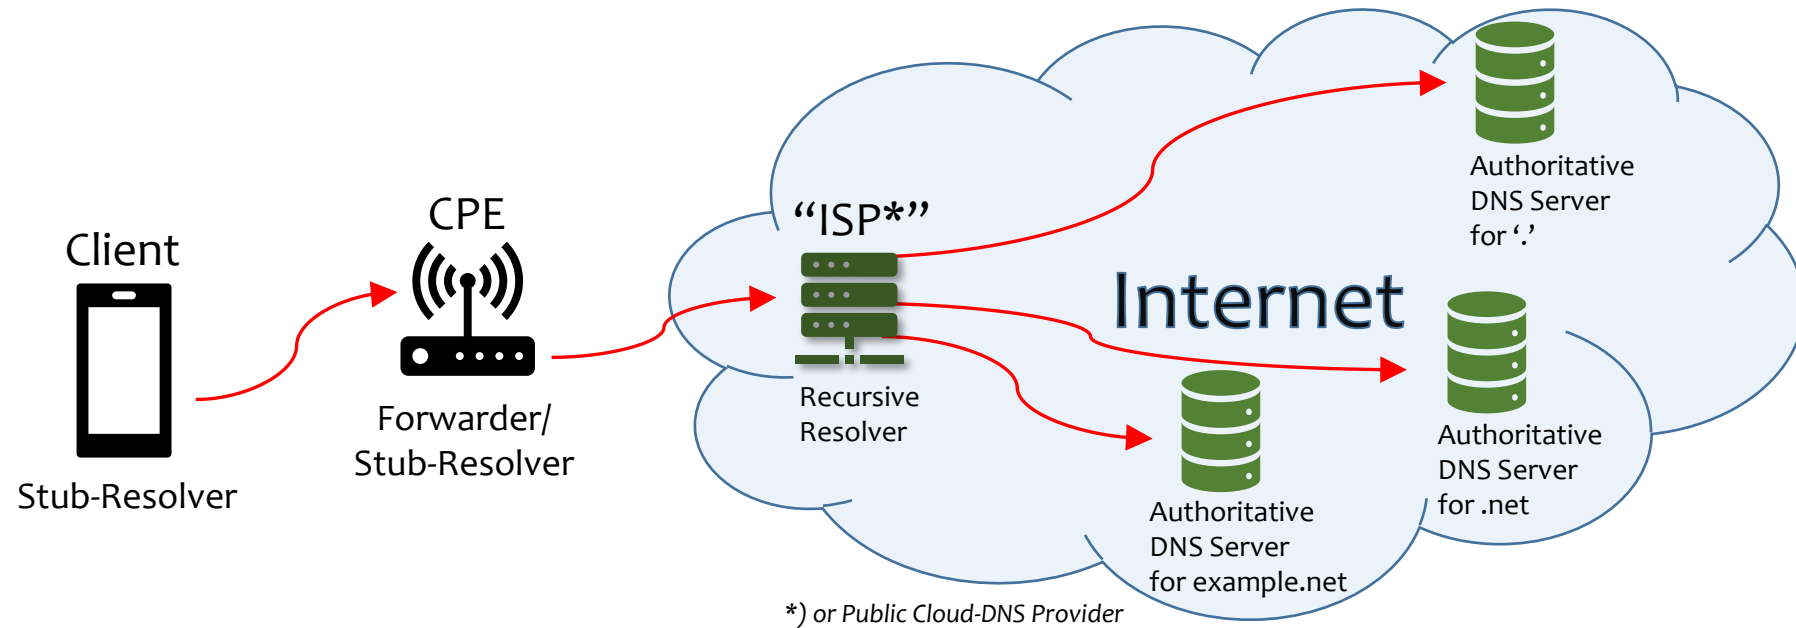
We compare Farsight's data  
to more than 1.1k zones as of August 2022

## Active Measurements

- Alexa Top 1 M,  
Aug 2022 (498k)
- One VP, 56 M queries
- Oct 11-14 & 22-24
- Developed own system:
  - workers
  - caches
  - DNS-probes
- Code:
  - partly PoC – level
  - written in ~ 3 months
  - but: working



# Stub-Resolvers and Recursion



No correlation between IP protocol used for DNS-resolution and protocol specific resource record types, i.e. A/AAAA



# A Short DNS Primer



- Zones, zone delegation
- Authoritative name servers
- In/out-of-bailiwick
- GLUE records
- Stub resolvers and recursion
- DNS using IPv4/IPv6 vs. A/AAAA resource records

