

# A look at authoritative nameserver performance of TLDs

**Wouter de Vries**  
Systems engineer

**DNS-OARC 43**

# Overview

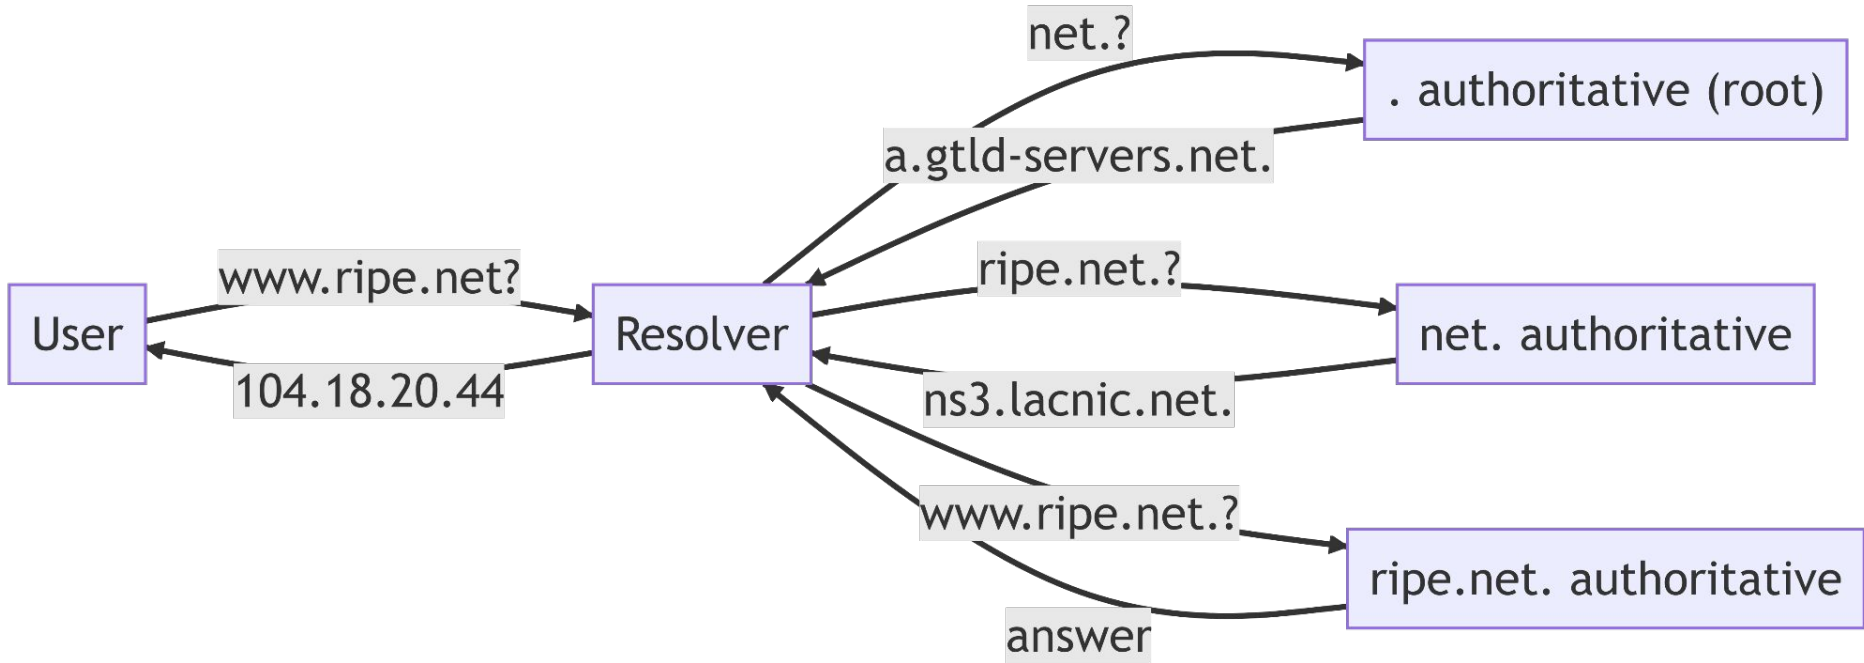
- Less studied than the root nameservers
- Higher impact on end-user QoE
- ~1500 different TLDs currently in the Root zone
- Each potentially having their own performance characteristics

# Overview

- Less studied than the root nameservers
- Higher impact on end-user QoE
- ~1500 different TLDs currently in the Root zone
- Each potentially having their own performance characteristics

# DNS? TLDs? Root?

## A quick recap

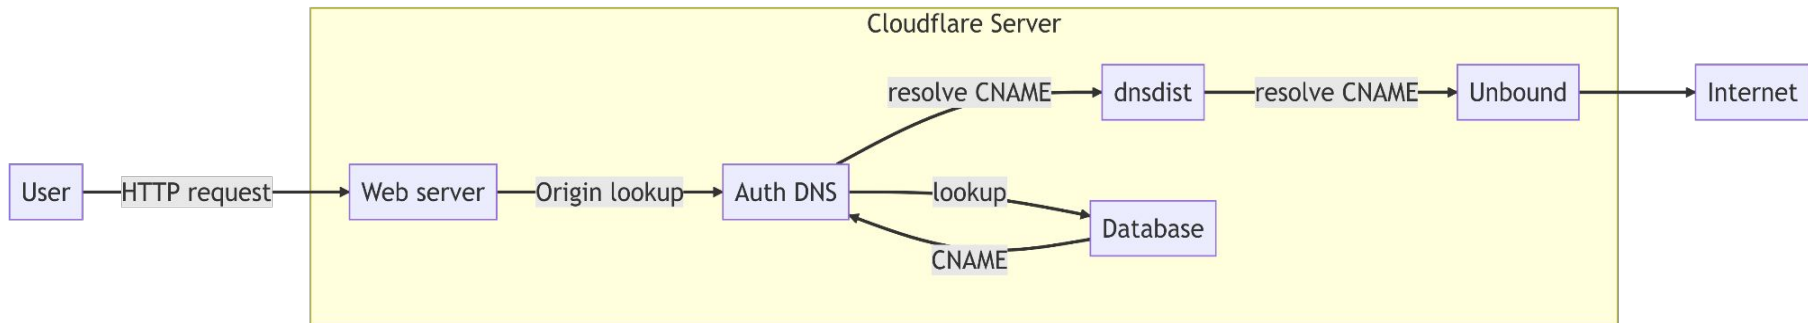


## What about Root?

- Latency is less important
- Low cardinality (~15k TLDs in total)
  - Easy to keep in cache
- TTLs are high, typically 24 or 48 hours.
  - **In theory** only a single lookup for .com per 48 hours

# When do we query?

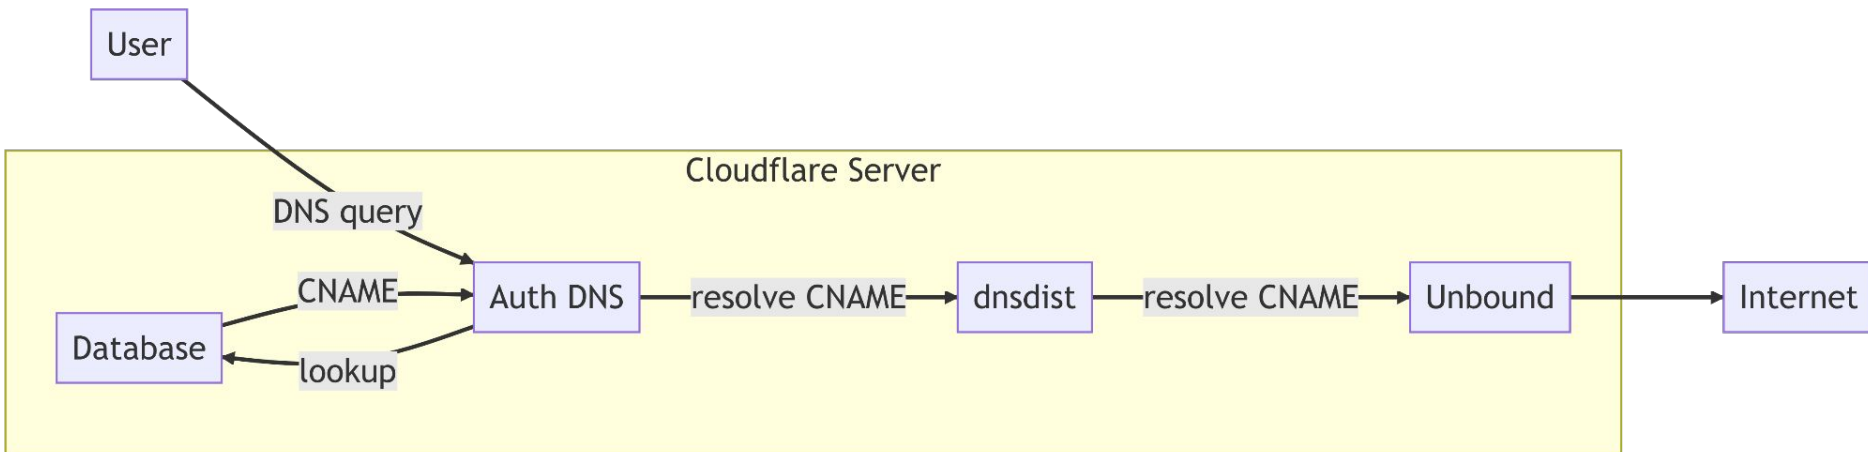
## *Authoritative DNS*



# When do we query?

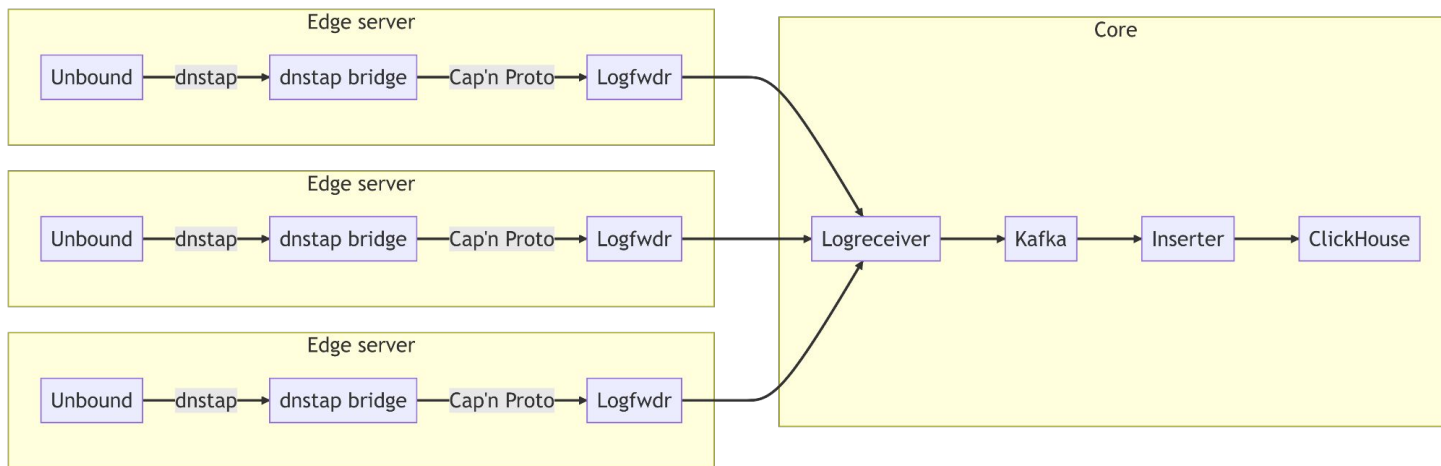
## *Authoritative DNS*

- DNS queries for CNAMEs that have *flattening* enabled



# Collecting data

- Collected at our edge servers in dnstap format
- Converted to Cap'n Proto for compatibility with internal services
- Stored in ClickHouse for analysis





# Collecting data

- Keep a daily updated dictionary in clickhouse
  - Key: TLD + IP address (e.g. ("nl", 194.0.25.24))
  - Value: nameserver (e.g. "ns3.dns.nl.")
- Materialized view using AggregatingMergeTree aggregating by:
  - Timeslot (1 hour granularity)
  - TLD
  - Datacenter ("colo")
  - Nameserver (e.g. "ns3.dns.nl.")
  - IPv4 (if v4)
  - IPv6 (if v6)
  - protocolType (UDP/TCP)
  - responseCode

# Collecting data

- Query counts
- Response time quantiles (0.25, 0.5, 0.75, 0.9, 0.95, 0.99, 0.999)
- Response time average

Currently ~5TiB of data  
collecting since 6th of March 2023

Total query count ~2.6 trillion

```
select count(*)  
from tld_metrics_1h
```

```
count()  
6291923456
```

```
select min(timeslot)  
from tld_metrics_1h
```

```
min(timeslot)  
2023-03-06 02:00:00
```


## Some numbers

- >3.5 billion queries to TLD authoritative nameservers per day
  - ~40k queries per second
- Roughly half to .com

# .com (and .net) latency

Median: **~7ms**

99th Percentile: **~183ms**



Winner of .com nameserver competition: [b.gtld-servers.net](https://b.gtld-servers.net)

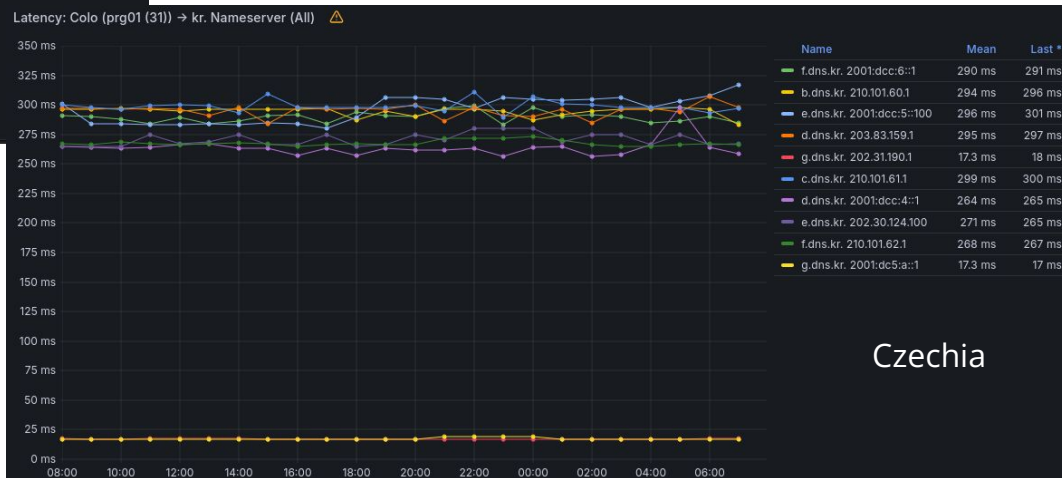
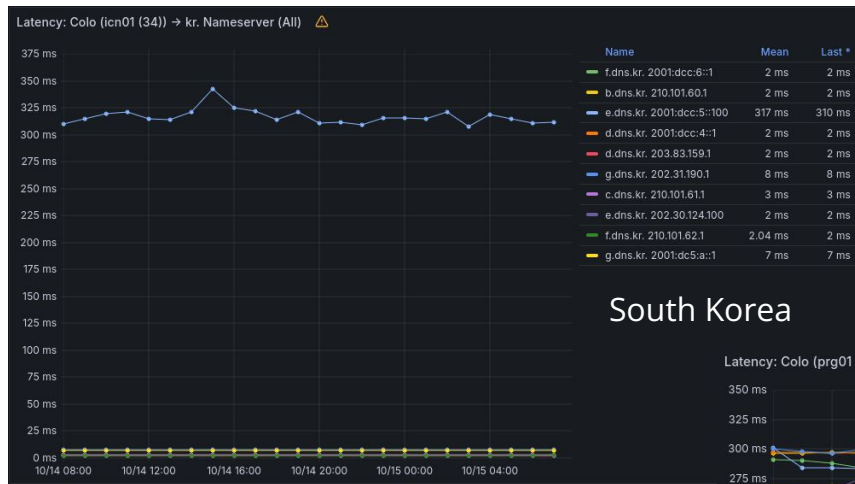
Median: **2.34ms** (IPv4), **6.3ms** (IPv6)

## Examples of slow TLDs

<b>TLD</b>	<b>p25</b>	<b>p50 (median)</b>	<b>p99</b>
kr.	47ms	122ms	346ms
cf.	40ms	101ms	359ms
ml.	40ms	98ms	358ms
top.	48ms	88ms	270ms
ir.	12ms	71ms	267ms

- Including only those TLDs we've sent at least a million queries to that day

# Examples of slow TLDs




## Examples of fast TLDs

<b>TLD</b>	<b>p25</b>	<b>p50 (median)</b>	<b>p99</b>
pl.	1ms	1ms	138ms
pt.	1ms	1ms	139ms
ca.	1ms	1ms	160ms
ar.	1ms	1ms	189ms
pw.	1ms	2ms	115ms

- Including only those TLDs we've sent at least a million queries to that day

# Examples of fast TLDs





## DNS-OARC 43 Host country ccTLD: .cz

- Median: **~16ms**
- 99th Percentile: **~167ms**
- Slightly faster in Seattle than in Prague :-)

nameserver	IP version	p25	p50 (median)	p99
b.ns.nic.cz.	IPv6	2ms	11ms	172ms
c.ns.nic.cz.	IPv4	1ms	13ms	158ms
c.ns.nic.cz.	IPv6	1ms	13ms	147ms
d.ns.nic.cz.	IPv6	3ms	14ms	156ms
a.ns.nic.cz.	IPv6	6ms	20ms	165ms
d.ns.nic.cz.	IPv4	1ms	22ms	177ms
a.ns.nic.cz.	IPv4	8ms	22ms	179ms
b.ns.nic.cz.	IPv4	1ms	24ms	210ms

# Current IP versus legacy IP

- Many (almost all!) TLDs support IPv6!
- 1457 -> yes, 24 -> no

region	ipVersion	p25	p50	p99
AFR	IPv4	7	50	299
AFR	IPv6	32	99	270
APAC	IPv4	2	29	214
APAC	IPv6	2	35	235
CHINA	IPv4	15	36	272
CHINA	IPv6	23	49	282
EEUR	IPv4	2	11	153
EEUR	IPv6	2	10	154
ENAM	IPv4	1	2	139
ENAM	IPv6	1	2	130
ME	IPv4	57	75	196
ME	IPv6	62	78	185
OC	IPv4	2	13	217
OC	IPv6	2	13	258
SAM	IPv4	8	50	225
SAM	IPv6	51	88	274
WEUR	IPv4	2	8	156
WEUR	IPv6	2	8	138
WNAM	IPv4	1	2	164
WNAM	IPv6	1	2	146

# How many IPs do you need?

- Big differences between TLDs
- The range is between 0 and 13 IP addresses per IP version

**.com, .net, .edu:** 13 IPv4, 13 IPv6

**.kp, .mh, .pf, .sl, .et, .sr:** 2 IPv4, 0 IPv6

# More IPs, More better?

- In terms of latency: not necessarily
  - **.nl:** 3 IPv4, 3 IPv6 - Median latency: 4ms, 7ms
  - **.net:** 13 IPv4, 13 IPv6 - Median latency: 7ms, 7ms
- Lots of variations
  - no clear correlation between number of IPs and performance

## More IPs == better?

- Anycast potentially hides a large network behind a few IPs
- Resolvers typically like faster IPs, but will keep querying slower IPs too
- Where are the users?
  - Does a ccTLD require a global anycast network? Or just a few well placed servers?

# Conclusions

## Folks picking a new domain:

- Consider the performance implications of the TLD
  - 100ms of added latency is noticeable, and that's when the actual processing of the request hasn't even started yet
- Not all TLDs are fast everywhere or slow everywhere

## TLD operators:

- Consider who is querying you
- What can be done to improve latency?
- Just adding more IPs doesn't necessarily improve performance (may in fact hurt performance)

# Questions