Passive analysis of DNS server reachability

14th CENTR R&D workshop

Maciej Andziński • maciej.andzinski@nic.cz • 29.05.2019

CZ.NIC | CZ DOMAIN REGISTRY

.CZ DNS servers

- 4 NS, all anycasted

```
cz.                 3600 IN   NS   a.ns.nic.cz.
cz.                 3600 IN   NS   b.ns.nic.cz.
cz.                 3600 IN   NS   c.ns.nic.cz.
cz.                 3600 IN   NS   d.ns.nic.cz.

a.ns.nic.cz.        3600 IN   A    194.0.12.1
a.ns.nic.cz.        3600 IN   AAAA 2001:678:f::1
b.ns.nic.cz.        3600 IN   A    194.0.13.1
b.ns.nic.cz.        3600 IN   AAAA 2001:678:10::1
d.ns.nic.cz.        3600 IN   A    193.29.206.1
d.ns.nic.cz.        3600 IN   AAAA 2001:678:1::1
```

Location of .CZ DNS servers

- **Asia**
  - [JP] Tokyo
- **Europe**
  - [AT] Vienna
  - [CZ] Undisclosed location, 2x Prague
  - [DE] Frankfurt
  - [SE] Stockholm
  - [UK] London
- **North America**
  - [US] California, Virginia
- **South America**
  - [BR] Sao Paulo
  - [CL] Santiago de Chile

12 locations

9 countries

4 continents

Motivation

- Help to answer the question:

  What is the best location for our DNS servers?

Motivation

- Help to answer the question:

  What is the best location for our DNS servers?

  - Who sends queries to our servers and
    how log does it take for a query to reach our server?

Motivation

- Help to answer the question:

  What is the best location for our DNS servers?

this is easy

  - Who sends queries to our servers and
    how log does it take for a query to reach our server?

this is challenging

CZ DOMAIN REGISTRY

Challenge

- How to measure the latency between a DNS client and a DNS server?
  - A typical solution: active measurements
    - PING from DNS server to DNS client
    - PING to DNS server from a probe (e.g. RIPE Atlas)

Our concept: passive analysis

- We capture DNS traffic that hits .CZ DNS servers
- There was **15,659,160,884** in the first two weeks of May 2019
  - UDP **15,638,098,643** queries **(99.87%)**
  - TCP: **21,062,241** queries **(0.13%)**

Our concept: passive analysis

- We capture DNS traffic that hits .CZ DNS servers
- There was **15,659,160,884** in the first two weeks of May 2019
  - UDP **15,638,098,643** queries **(99.87%)**
  - TCP: **21,062,241** queries **(0.13%)**

~ 17 TCP connections per second

- Let‘s use TCP data to evaluate the latency between a DNS client and a DNS server!

cz.nic | CZ DOMAIN REGISTRY

TCP handshake

TCP handshake

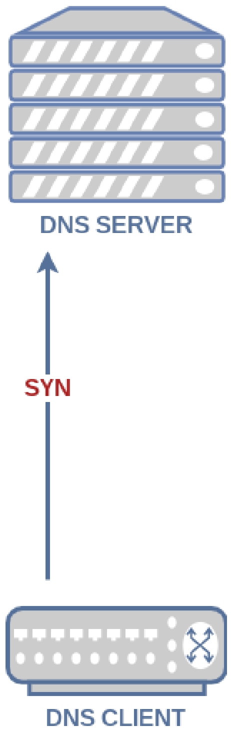

TCP handshake

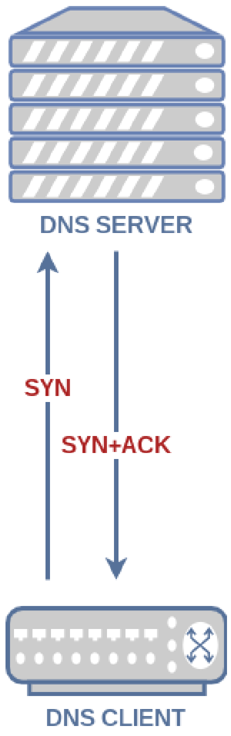

CZ.NIC | CZ DOMAIN REGISTRY

TCP handshake

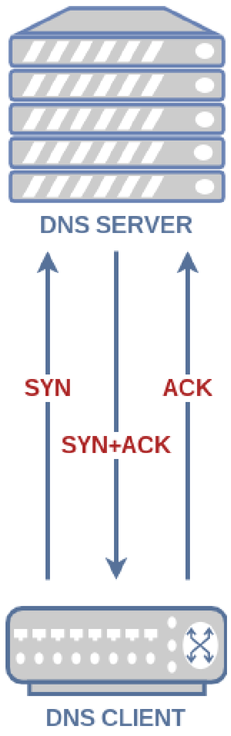

RTT of a TCP handshake

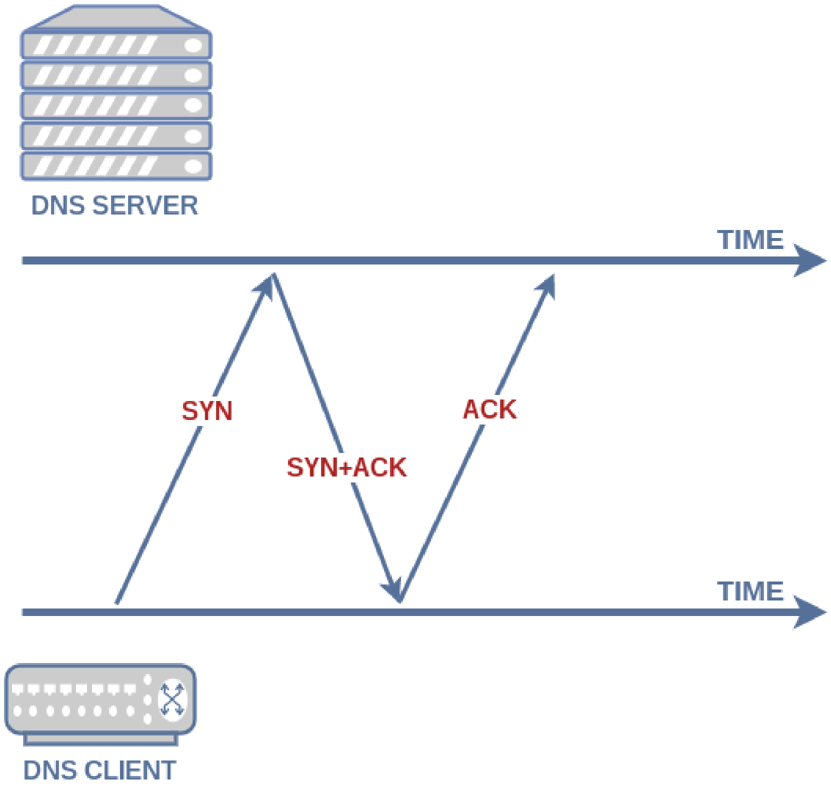

RTT of a TCP handshake

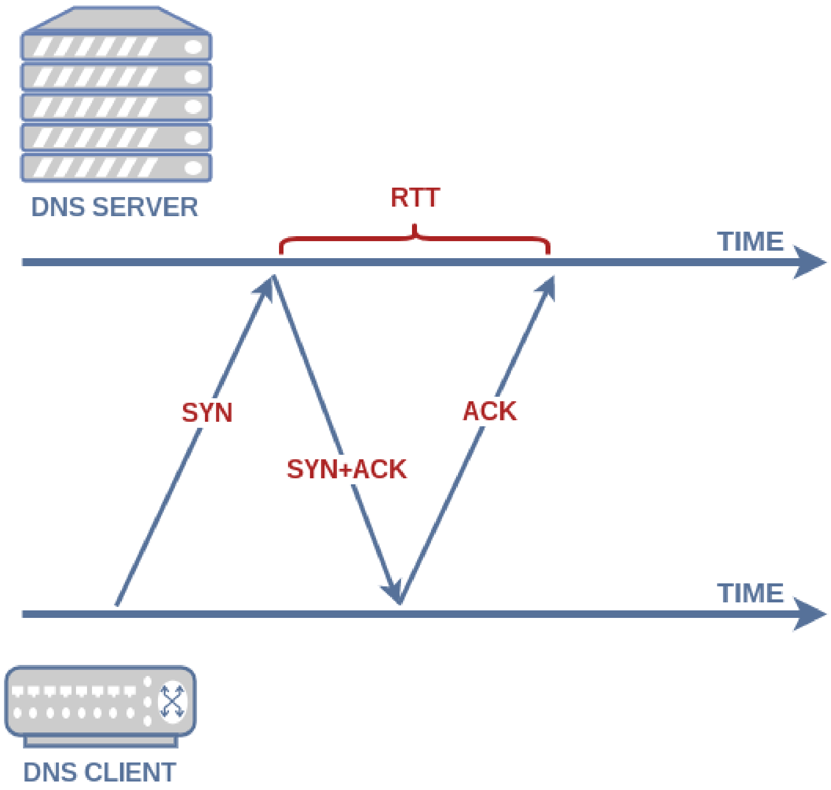

Our concept

1) For each pair (client, server) compute median RTT of a TCP handshake

client_ip	client_cc	client_asn	server	**queries**	tcp	**median_rtt**
217.31.193.164	CZ	25192	[Europe] AT, Vienna	**37123**	0	**NA**
217.31.193.164	CZ	25192	[Europe] CZ, Undisclosed	**5171434**	57	**12.7 ms**
217.31.193.164	CZ	25192	[Europe] CZ, Praha – CECOLO	**2579707**	6	**11.9 ms**
217.31.193.164	CZ	25192	[Europe] CZ, Praha - CRA	**27065563**	220	**11.5 ms**
217.31.193.164	CZ	25192	[Europe] UK, London	**8416765**	88	**43.4 ms**

Total number of DNS queries (UDP+TCP)

Number of captured TCP sessions

cz.nic | CZ DOMAIN REGISTRY

Our concept

2)Evaluate RTT for each client, network, country, …

(**Evaluated RTT** = *weighted mean* of RTT for all servers)

client_ip	client_cc	client_asn	server	queries	**median_rtt**	**weight**
217.31.193.164	CZ	25192	[Europe] AT, Vienna	37123	**NA**	**0.000858**
217.31.193.164	CZ	25192	[Europe] CZ, Undisclosed	5171434	**12.7 ms**	**0.120**
217.31.193.164	CZ	25192	[Europe] CZ, Praha – CECOLO	2579707	**11.9 ms**	**0.0596**
217.31.193.164	CZ	25192	[Europe] CZ, Praha - CRA	27065563	**11.5 ms**	**0.625**
217.31.193.164	CZ	25192	[Europe] UK, London	8416765	**43.4 ms**	**0.195**

$$RTT = \sum_{i=1}^{n} Norm(w_i) \cdot RTT_i \qquad for\ RTT_i \neq NA$$

Evaluated RTT for 217.31.193.164 = **17.9 ms**

CZ.NIC | CZ DOMAIN REGISTRY

Results

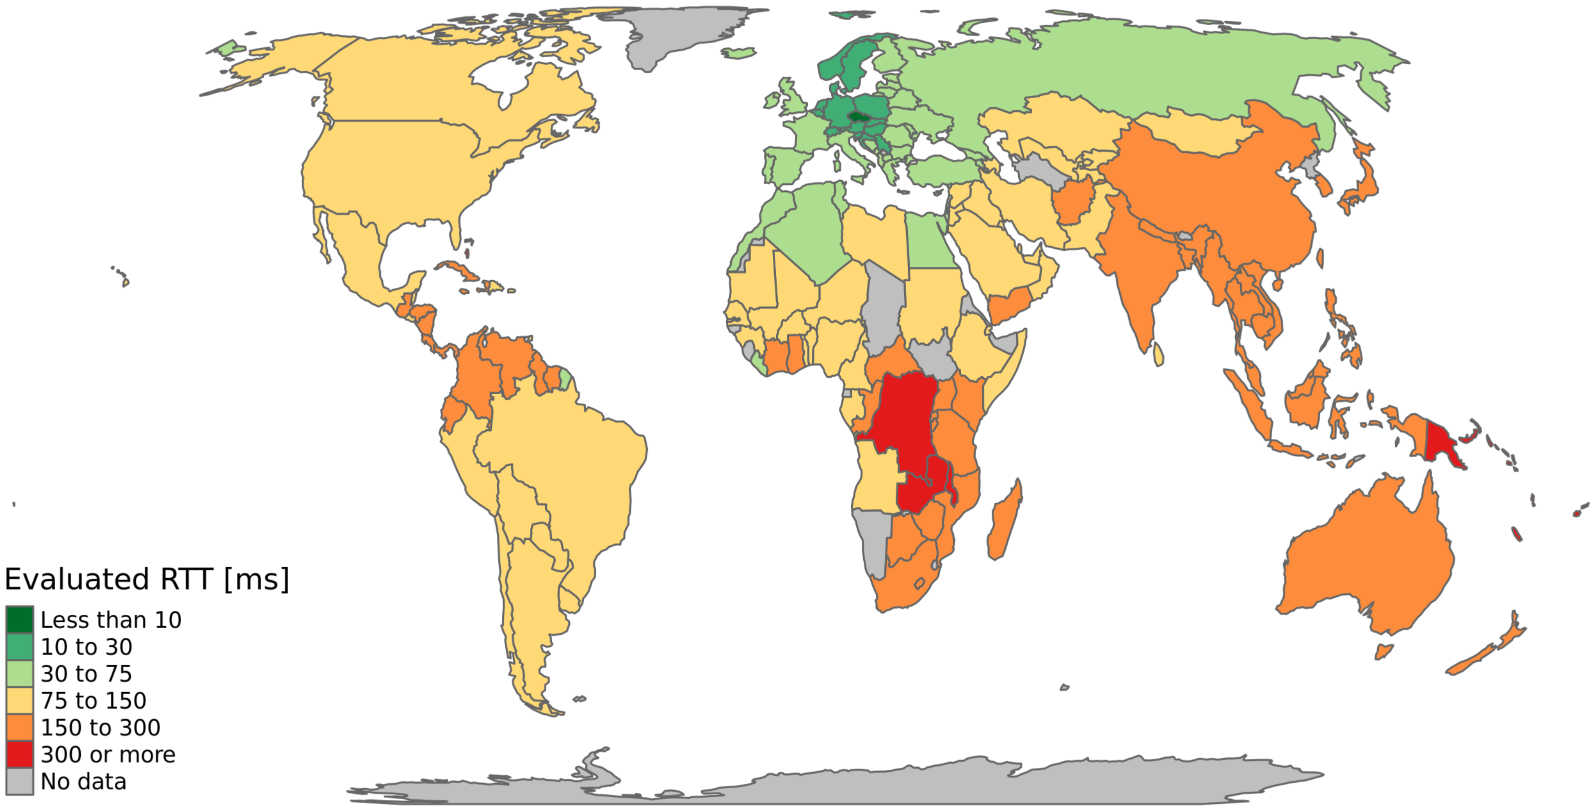

CZ DOMAIN REGISTRY

Results

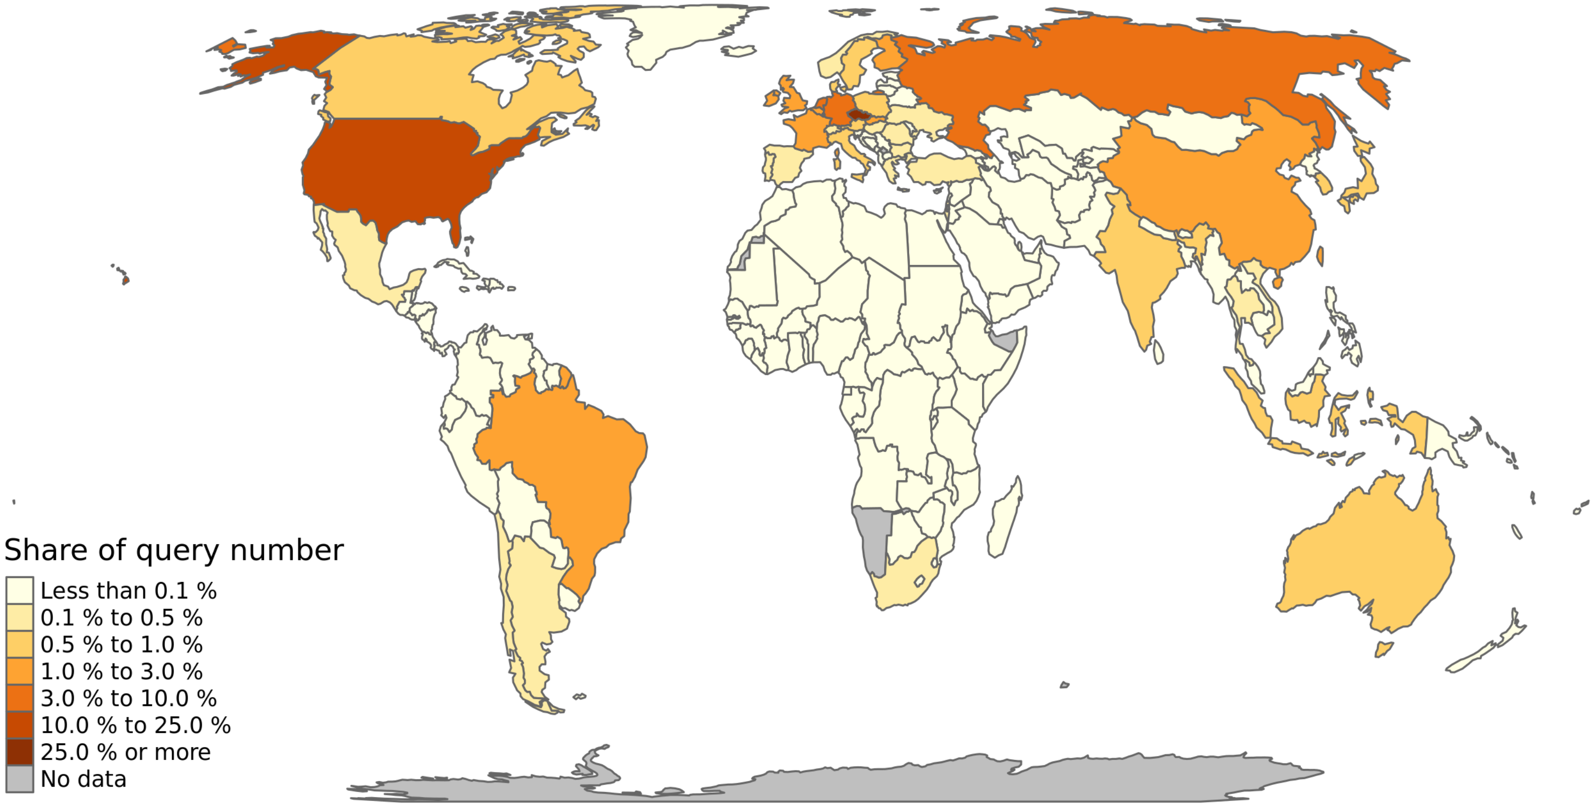

CZ DOMAIN REGISTRY

Results

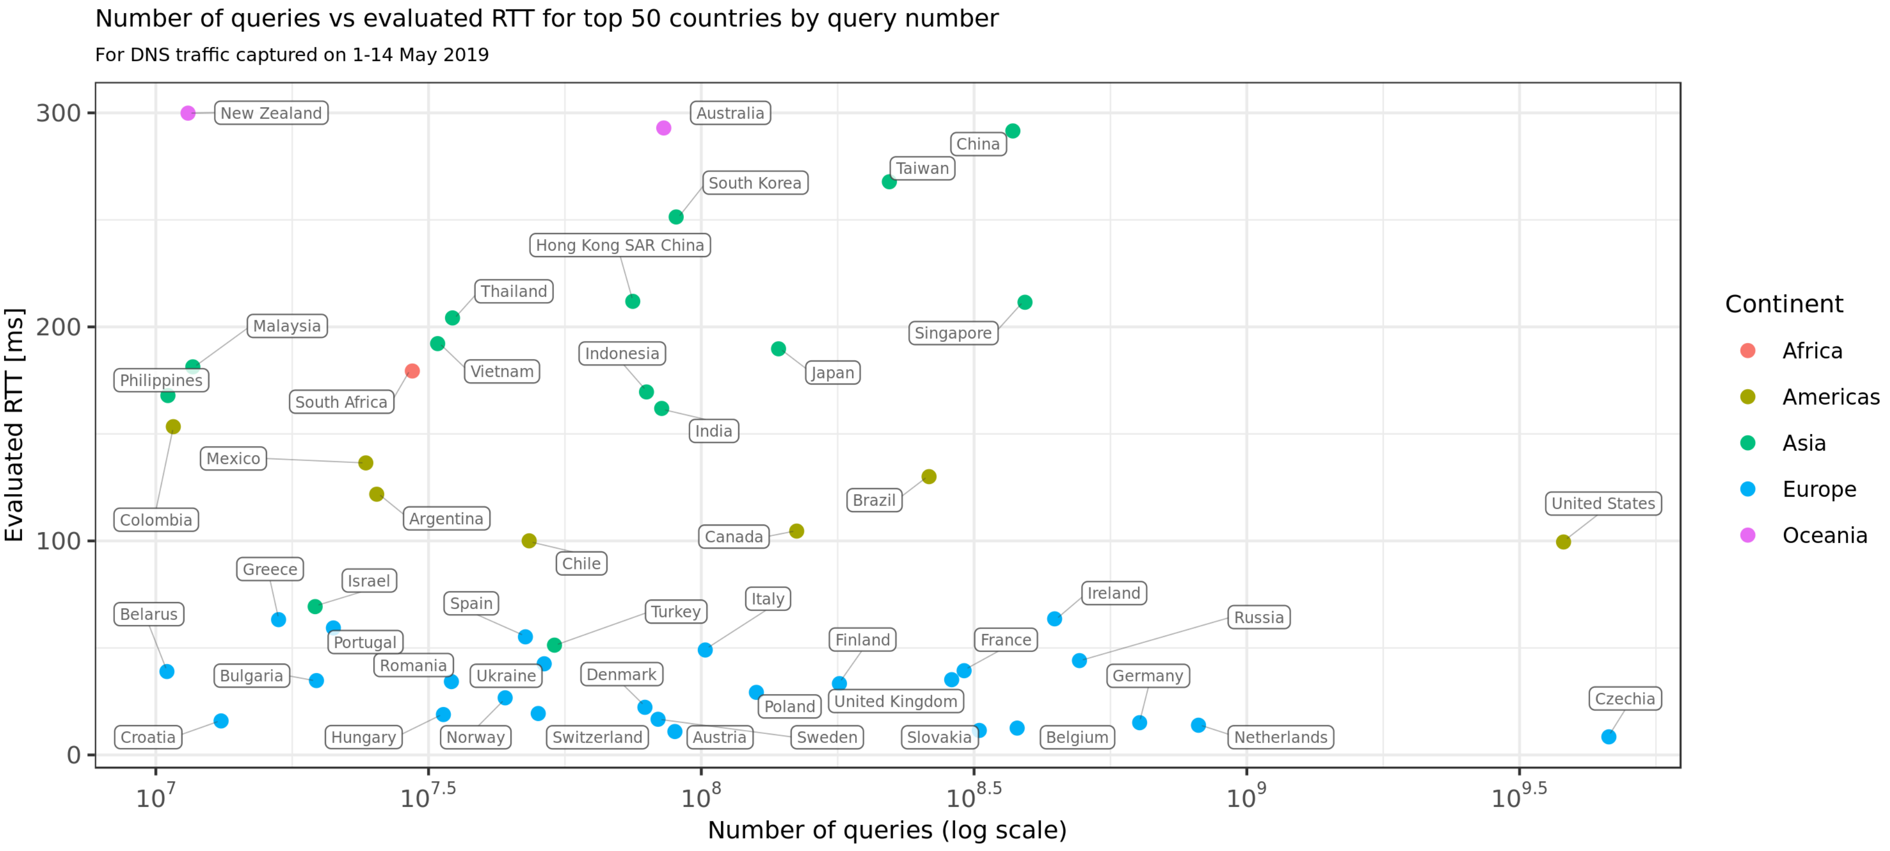

Number of queries vs evaluated RTT for top 50 countries by query number
For DNS traffic captured on 1-14 May 2019
Evaluated RTT [ms]
300
200
100
0
Number of queries (log scale)
10^7
$10^{7.5}$
10^8
$10^{8.5}$
10^9
$10^{9.5}$
New Zealand
Australia
China
Taiwan
South Korea
Hong Kong SAR China
Thailand
Singapore
Malaysia
Indonesia
Vietnam
Japan
Philippines
South Africa
India
Mexico
Brazil
Argentina
Colombia
Canada
United States
Chile
Greece
Israel
Belarus
Spain
Turkey
Italy
Ireland
Portugal
Russia
Finland
France
Romania
Bulgaria
Ukraine
Denmark
Germany
Poland
United Kingdom
Czechia
Croatia
Hungary
Norway
Switzerland
Austria
Sweden
Slovakia
Belgium
Netherlands
Continent
Africa
Americas
Asia
Europe
Oceania

Results

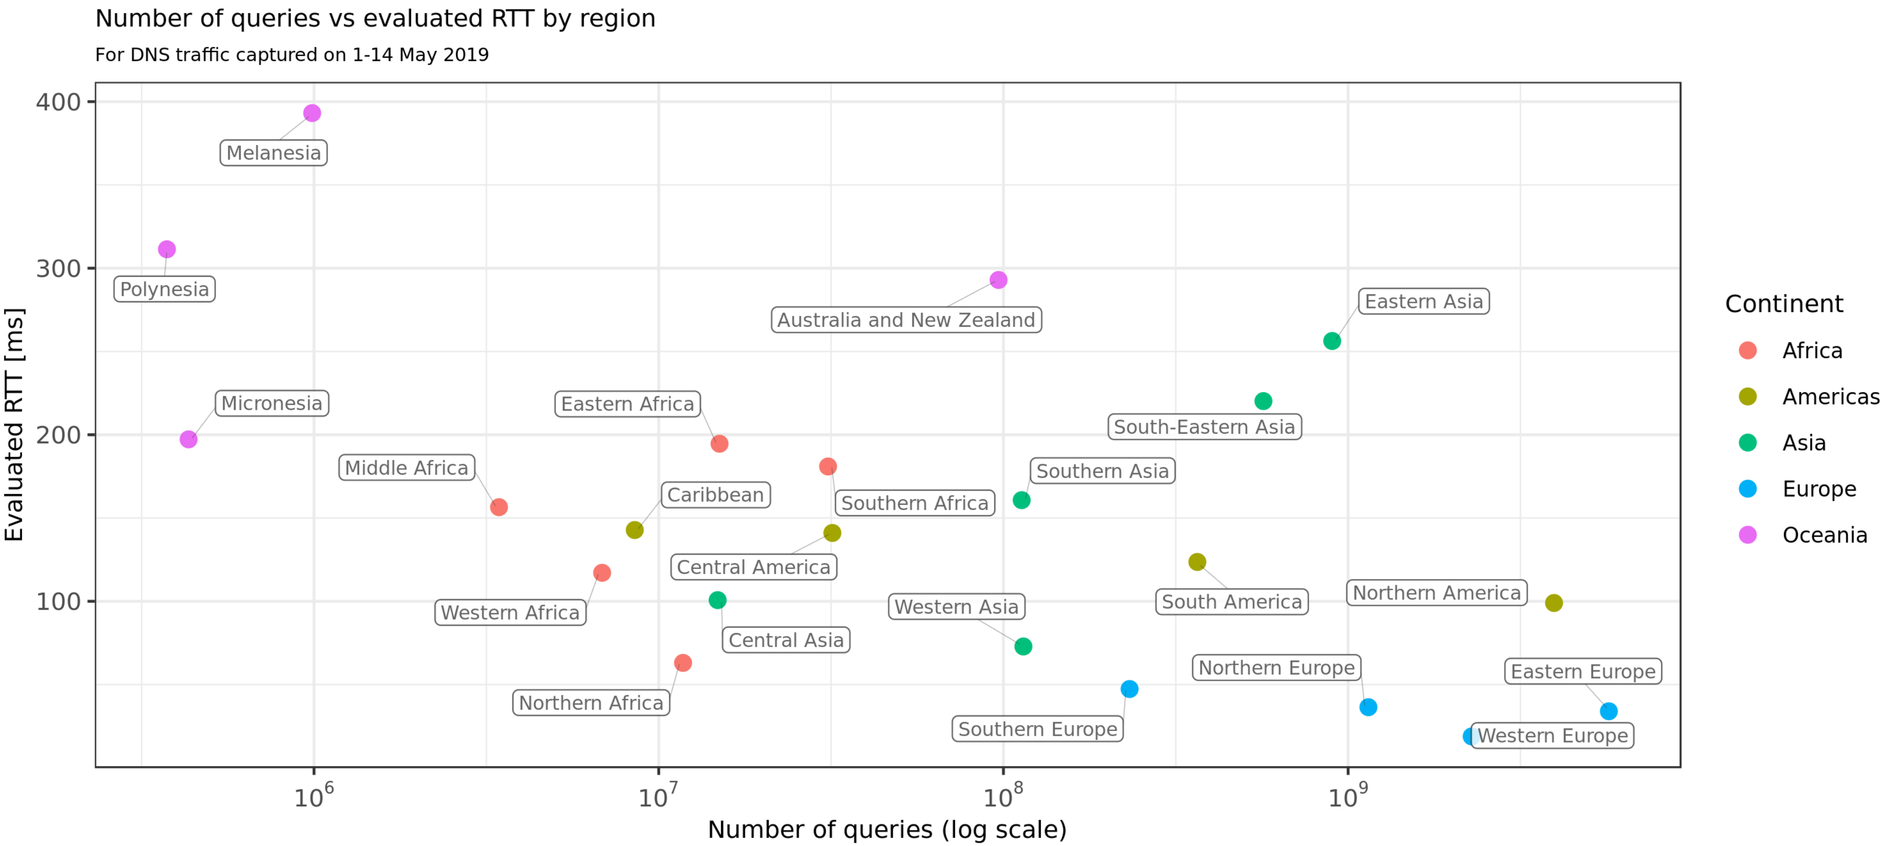

CZ.NIC CZ DOMAIN REGISTRY

Results

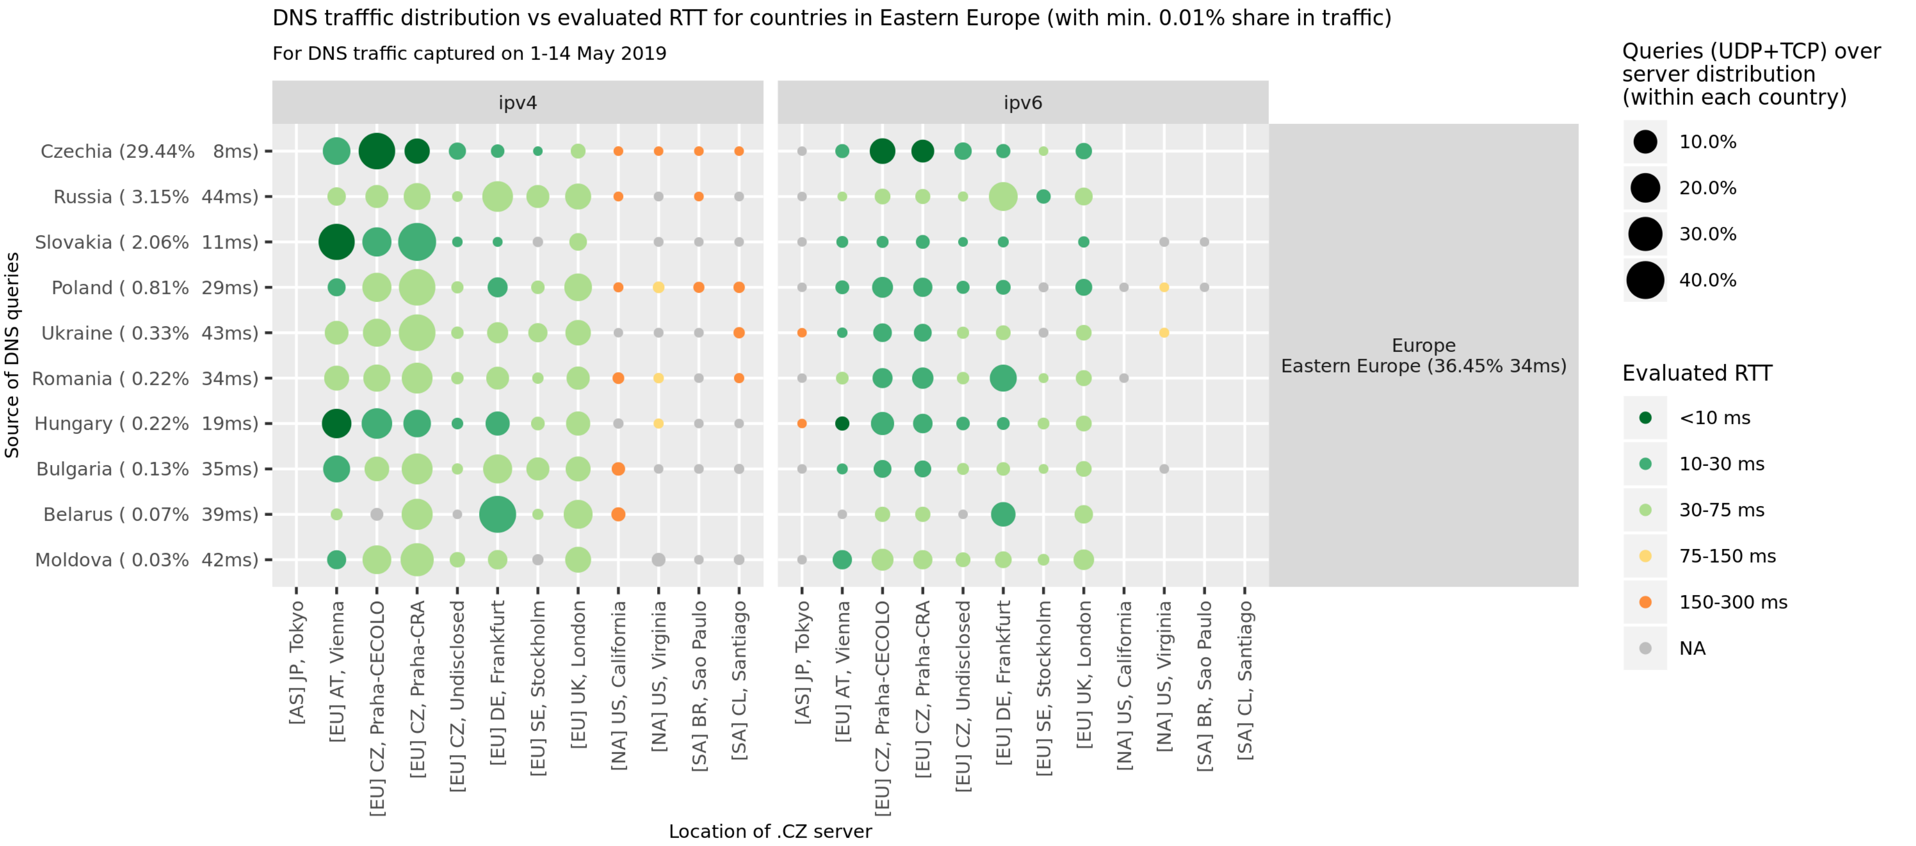

CZ DOMAIN REGISTRY

Results

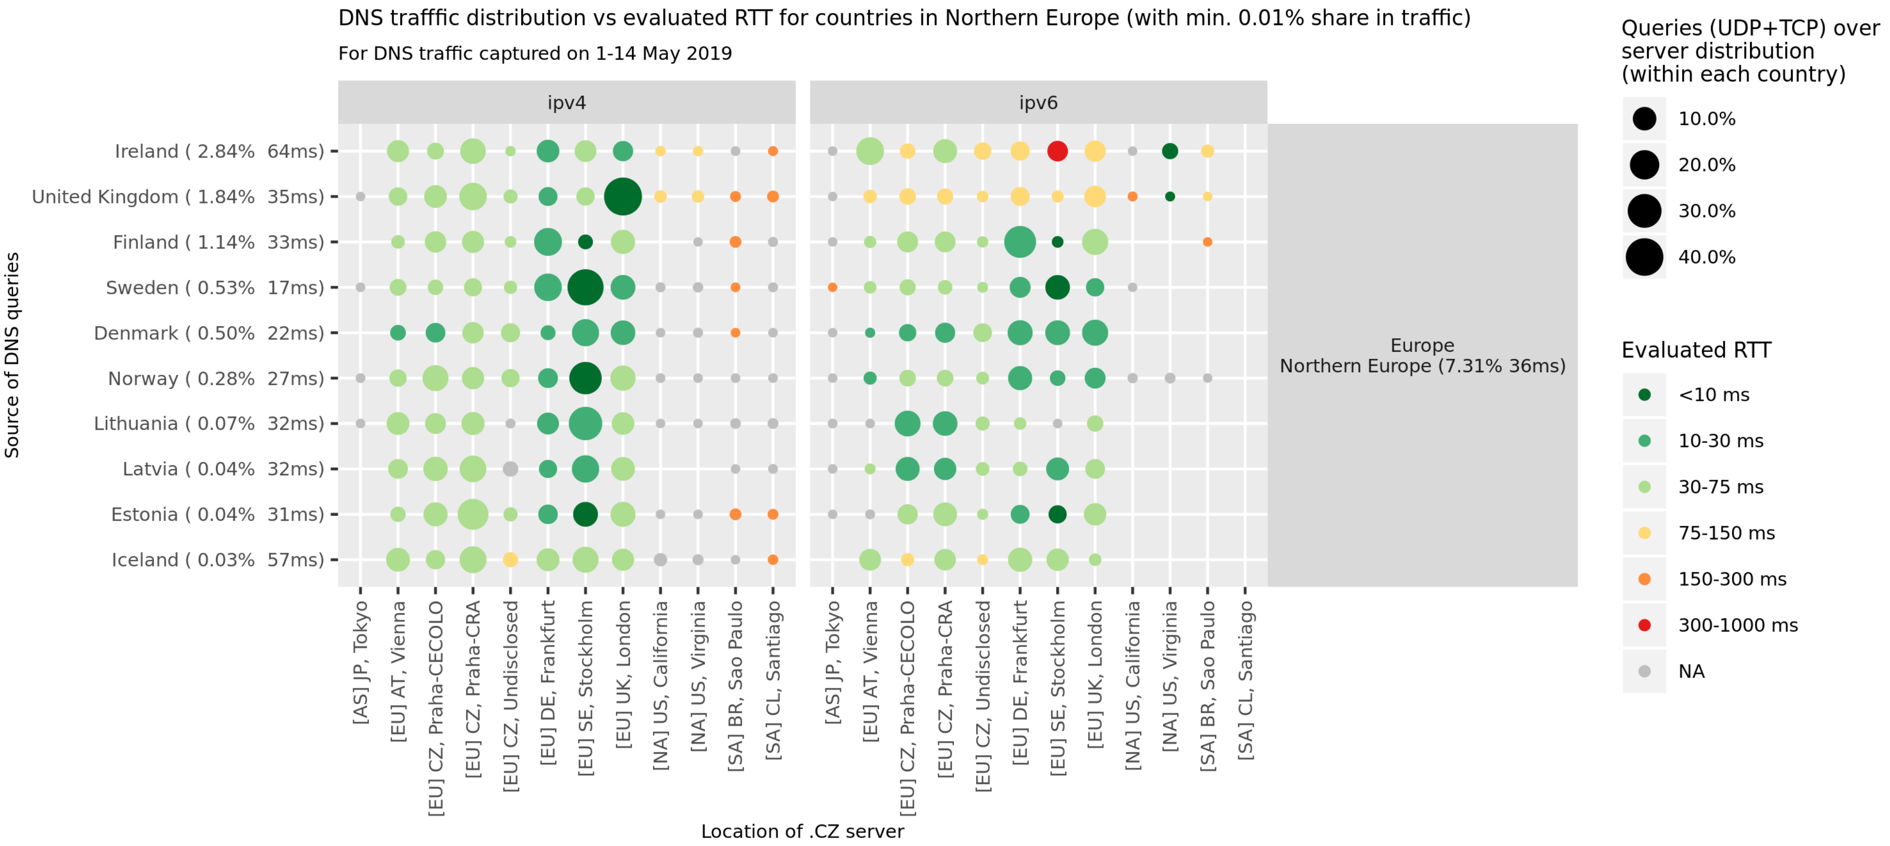

cz.nic CZ DOMAIN REGISTRY

Results

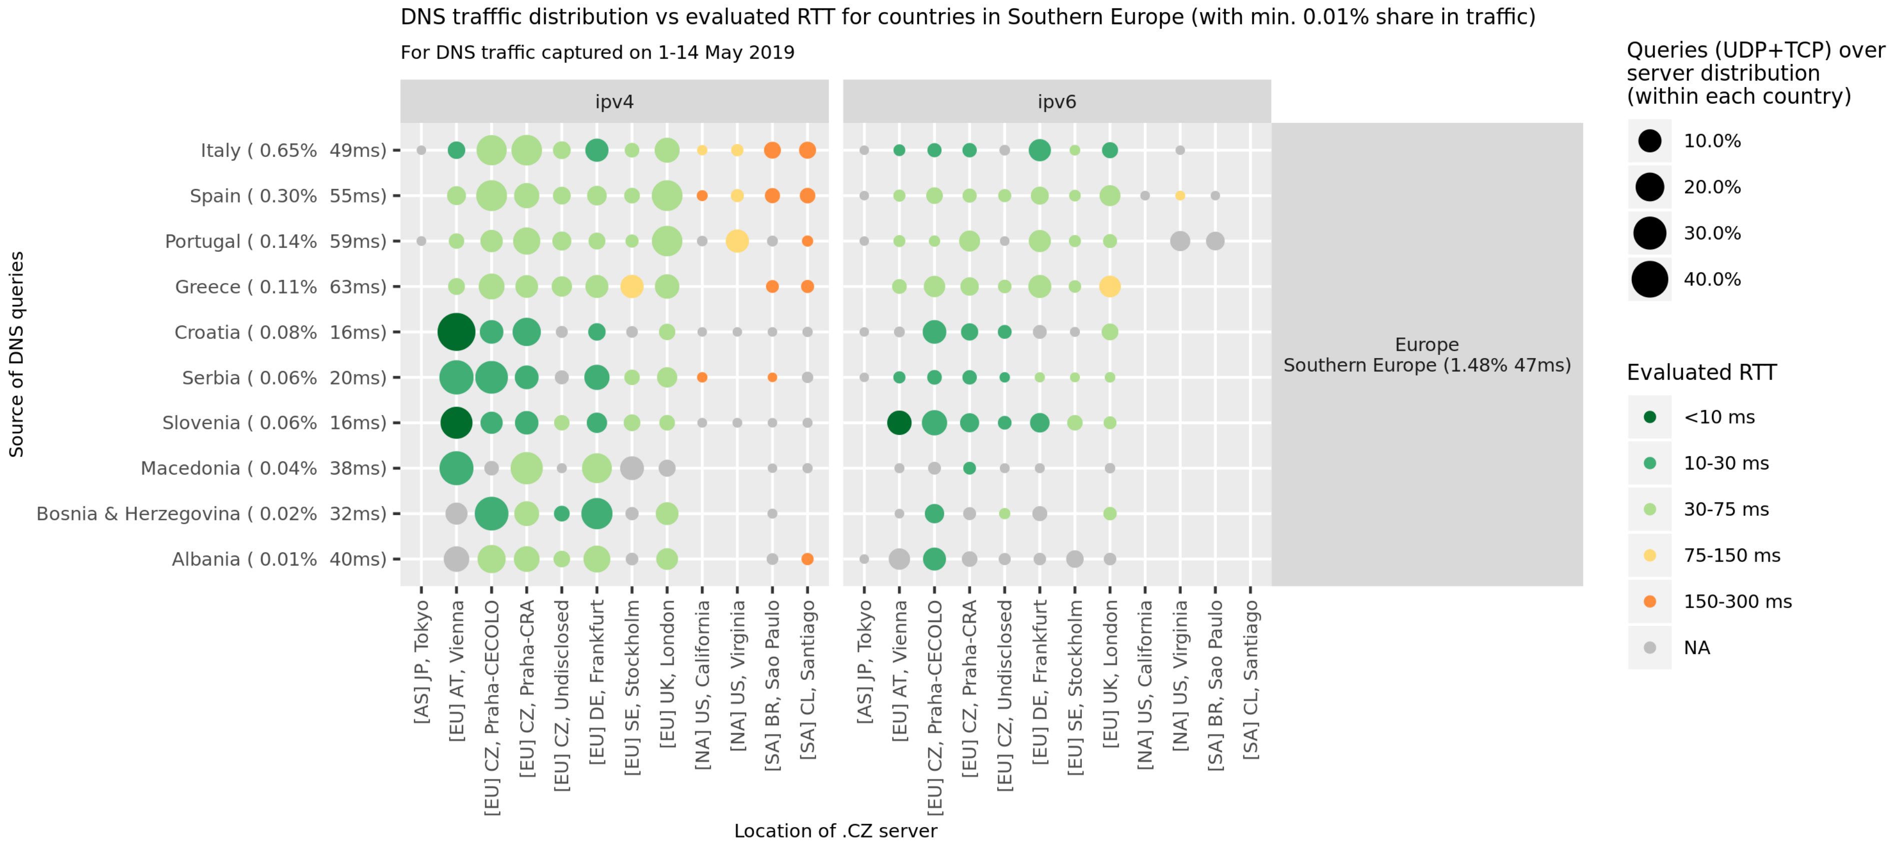

DNS trafffic distribution vs evaluated RTT for countries in Southern Europe (with min. 0.01% share in traffic)
For DNS traffic captured on 1-14 May 2019
ipv4
ipv6
Italy (0.65% 49ms)
Spain (0.30% 55ms)
Portugal (0.14% 59ms)
Greece (0.11% 63ms)
Croatia (0.08% 16ms)
Serbia (0.06% 20ms)
Slovenia (0.06% 16ms)
Macedonia (0.04% 38ms)
Bosnia & Herzegovina (0.02% 32ms)
Albania (0.01% 40ms)
Source of DNS queries
Europe
Southern Europe (1.48% 47ms)
[AS] JP, Tokyo
[EU] AT, Vienna
[EU] CZ, Praha-CECOLO
[EU] CZ, Praha-CRA
[EU] CZ, Undisclosed
[EU] DE, Frankfurt
[EU] SE, Stockholm
[EU] UK, London
[NA] US, California
[NA] US, Virginia
[SA] BR, Sao Paulo
[SA] CL, Santiago
Location of .CZ server
Queries (UDP+TCP) over server distribution (within each country)
10.0%
20.0%
30.0%
40.0%
Evaluated RTT
<10 ms
10-30 ms
30-75 ms
75-150 ms
150-300 ms
NA

Results

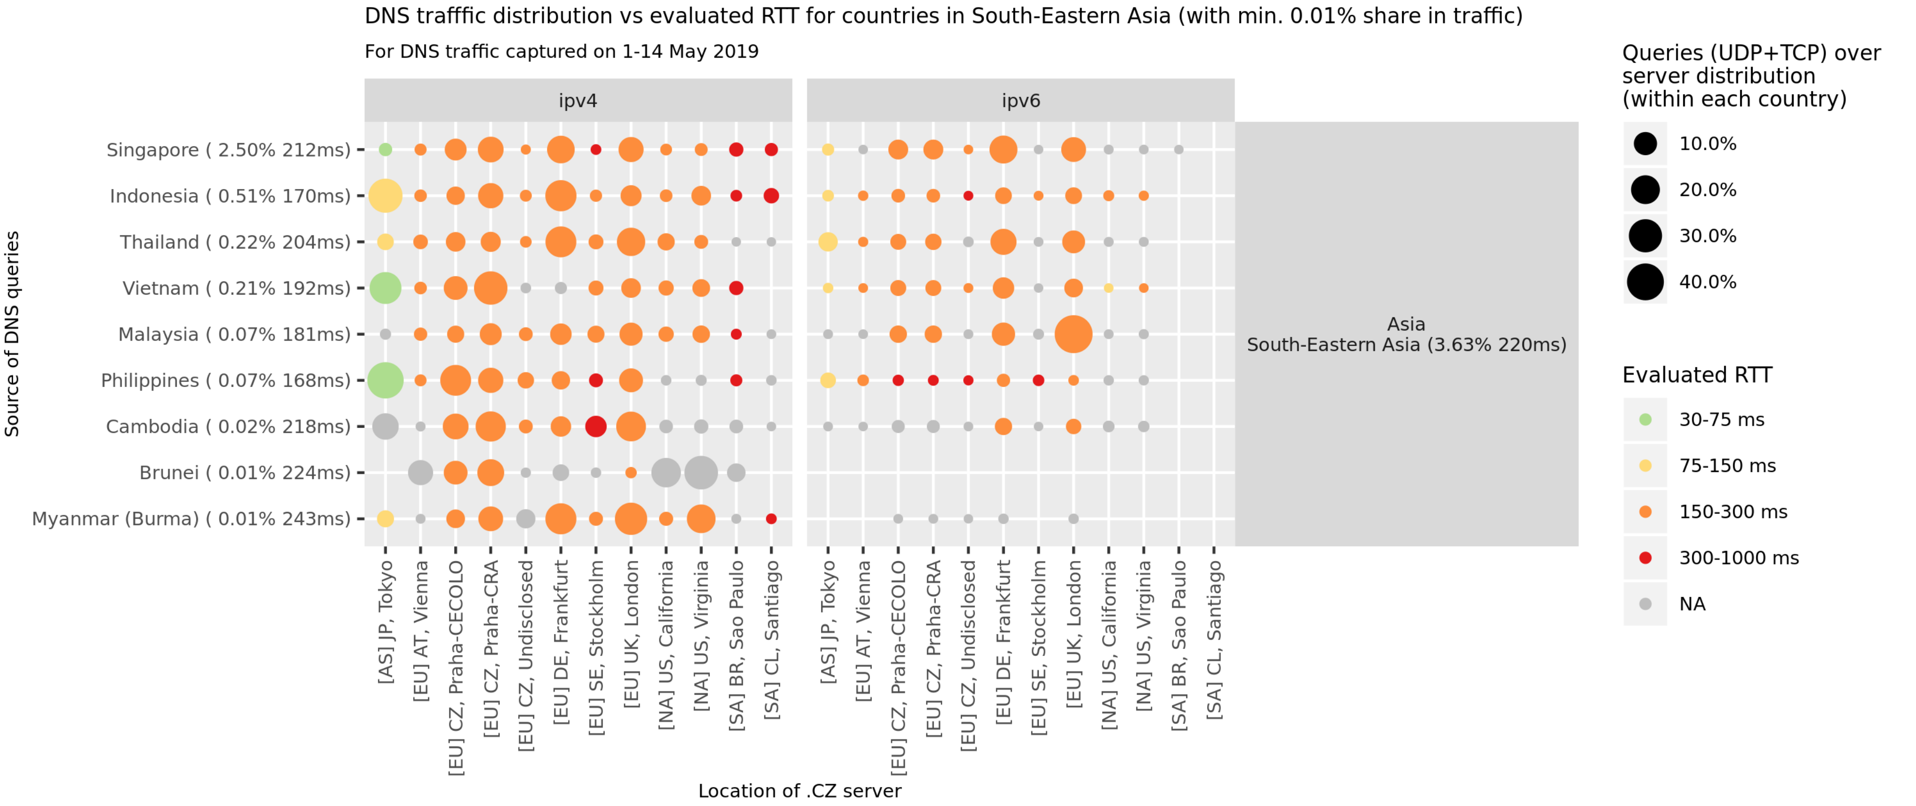

CZ DOMAIN REGISTRY

Results

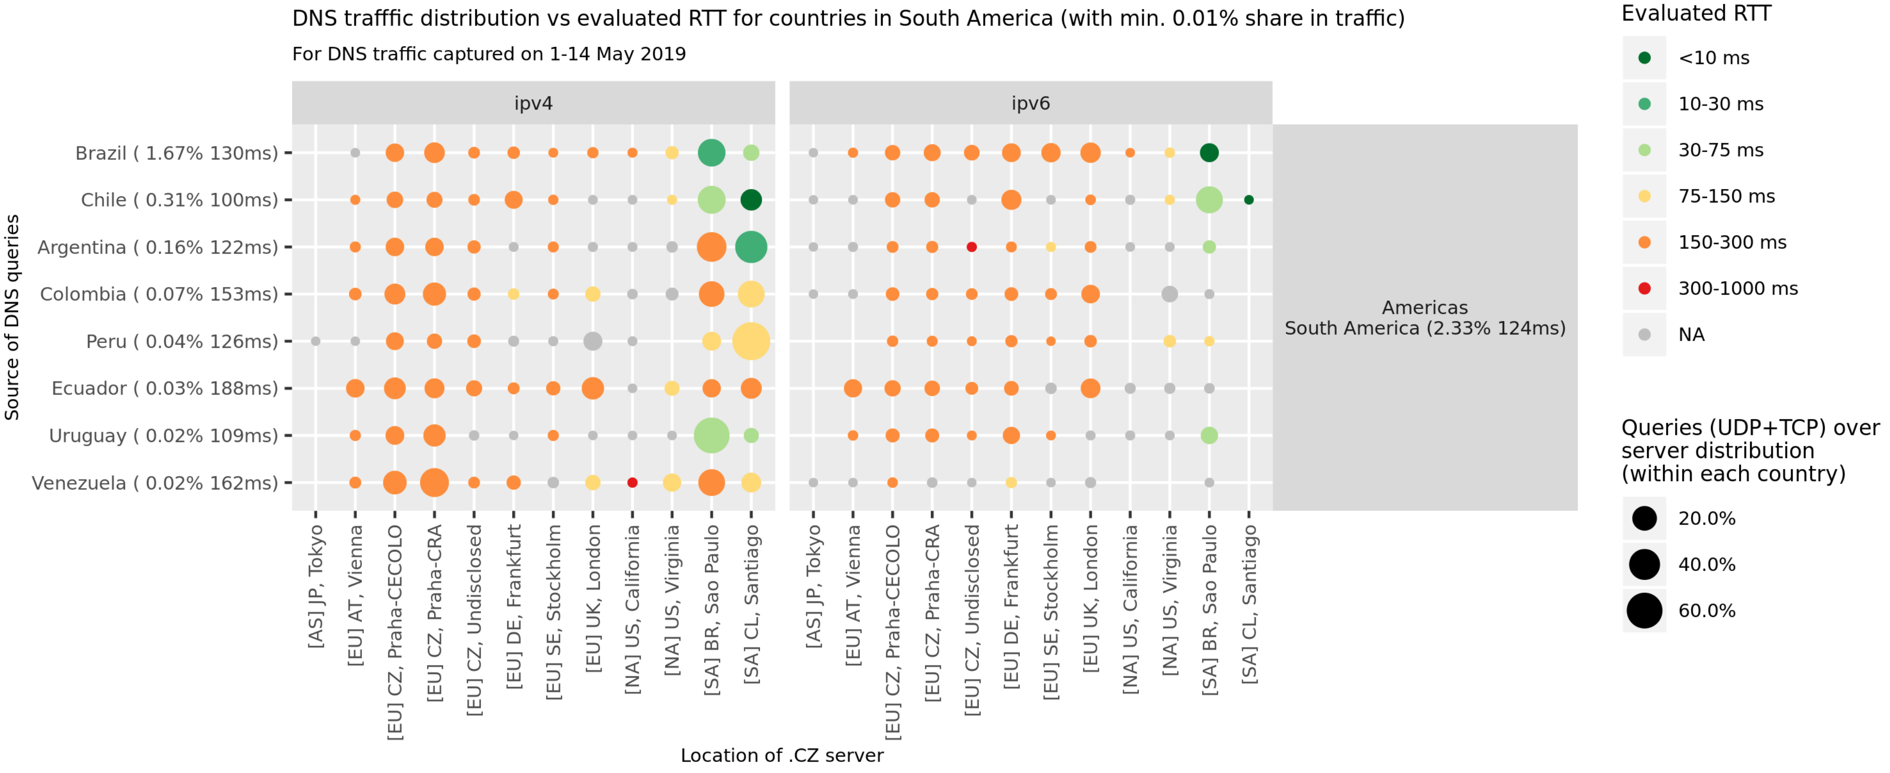

CZ DOMAIN REGISTRY

Results

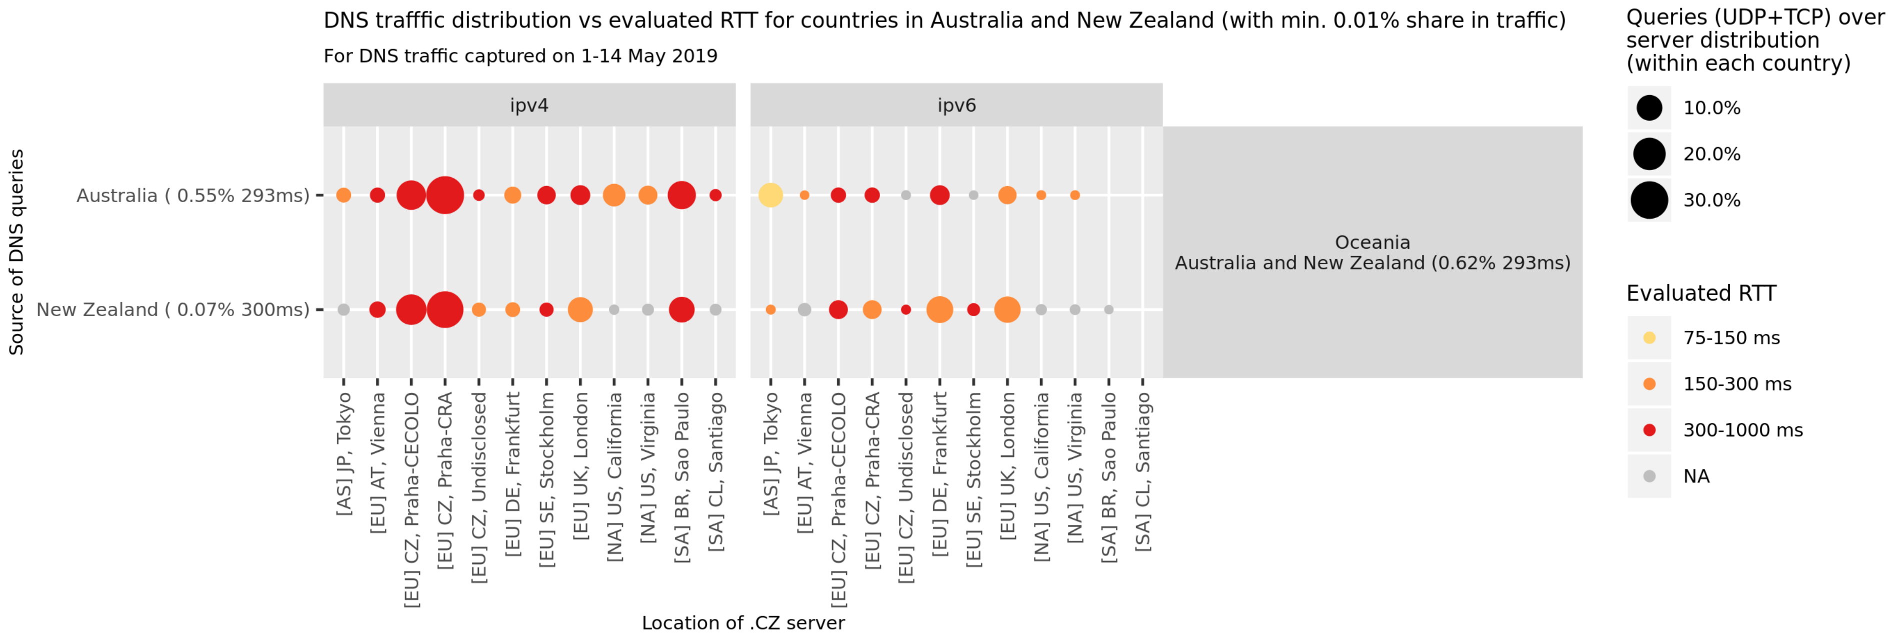
DNS trafffic distribution vs evaluated RTT for countries in Australia and New Zealand (with min. 0.01% share in traffic)
For DNS traffic captured on 1-14 May 2019
ipv4
ipv6
Source of DNS queries
Australia (0.55% 293ms)
New Zealand (0.07% 300ms)
Oceania
Australia and New Zealand (0.62% 293ms)
[AS] JP, Tokyo
[EU] AT, Vienna
[EU] CZ, Praha-CECOLO
[EU] CZ, Praha-CRA
[EU] CZ, Undisclosed
[EU] DE, Frankfurt
[EU] SE, Stockholm
[EU] UK, London
[NA] US, California
[NA] US, Virginia
[SA] BR, Sao Paulo
[SA] CL, Santiago
Location of .CZ server
Queries (UDP+TCP) over server distribution (within each country)
10.0%
20.0%
30.0%
Evaluated RTT
75-150 ms
150-300 ms
300-1000 ms
NA

Conclusion (on results)

- Geography matters
- Peering matters
- More than 1 server in a region needed to provide good RTT
  - Fewer NS → better?
- RTT: excellent in Czech Republic and very good in Europe (most of the traffic), but poor in some remote areas
  - A server down under may be a good idea

Conclusion (on method)

(-) Drawbacks

- Much traffic needed
- Sometimes difficult to measure RTT of TCP handshake (retransmissions, broken handshakes, lost packets)

(+) Advantages

- Delivers RTT for actual origin of a DNS query
- Relatively easy to deploy

(?) Other remarks

- Considerations on TCP occurence in DNS
- GeoIP accuracy / updates

CZ DOMAIN REGISTRY

Thank You

Maciej Andziński • maciej.andzinski@nic.cz

CZ.NIC | CZ DOMAIN REGISTRY