DDoS prevention in .cz

Tech Day – ICANN 66

Ondřej Filip • ondrej.filip@nic.cz • 4 Nov 2019 • Montreal

Trigger

- DDoS attack against domain .tr
  - Presented at Tech Day ICANN 60
  - “One ISP reported 220 Gbps attack bandwidth”
- Evaluated on our side - current anycast setup not sufficient
- Some HW tools evaluated – not preventing all attacks
- Result: **Anycast capacity upgrade – focus on Czechia**

Our anycast - World

Our anycast - Europe

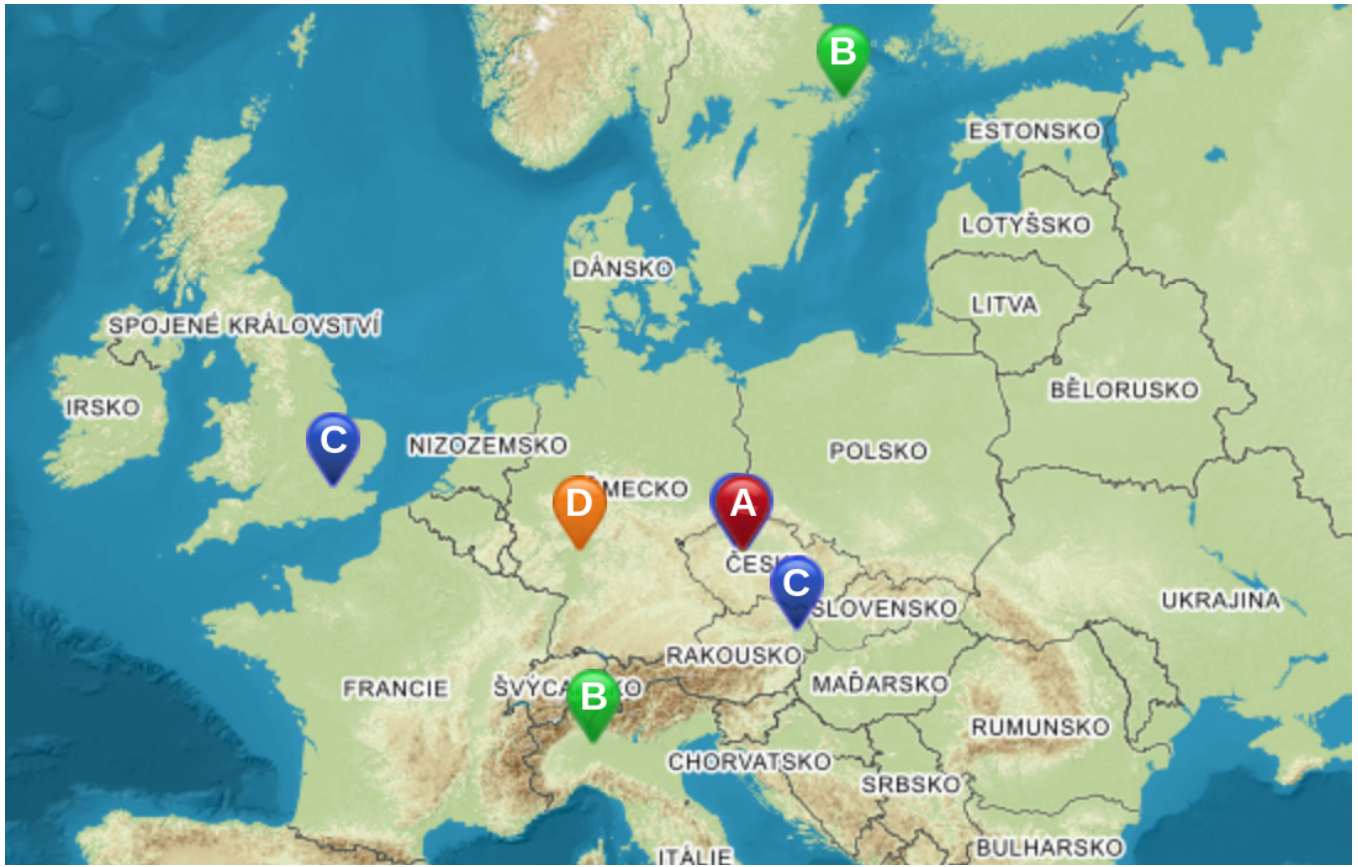

Our anycast - Czechia/Prague

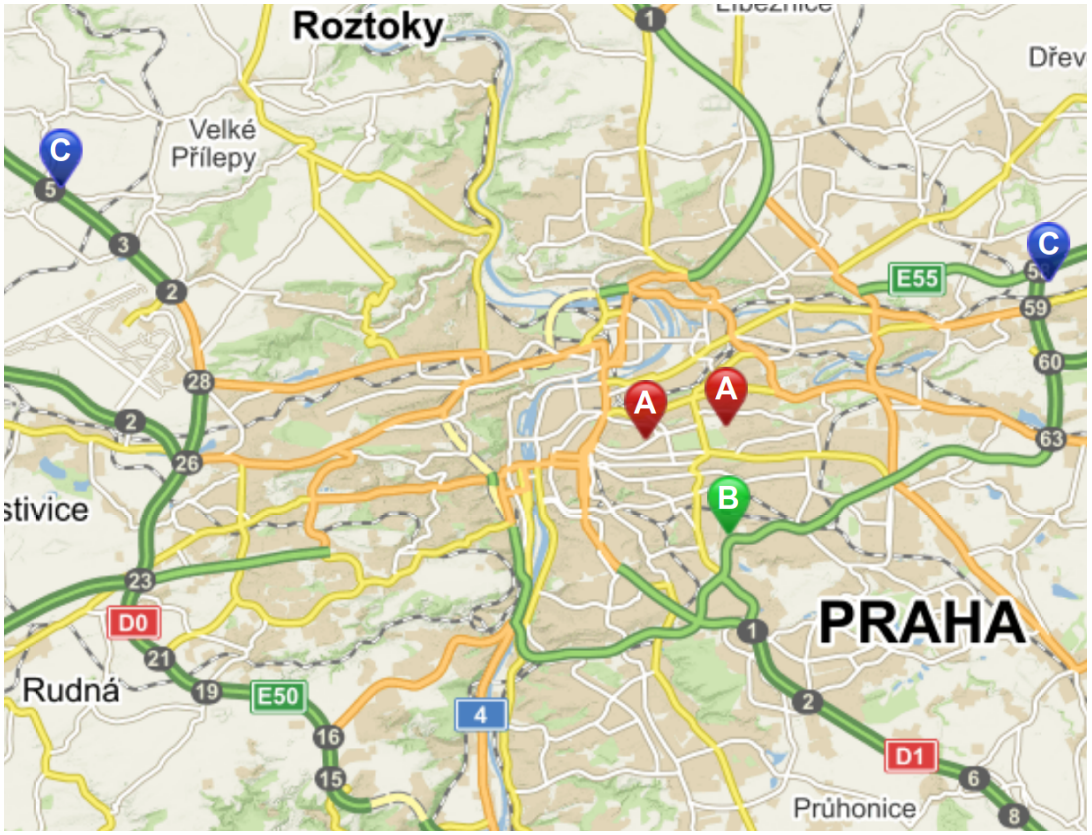

Basics for .CZ DNS anycast

- **Asia**
  - [JP] Tokyo
- **Europe**
  - [AT] Vienna
  - [CZ] 5 x Prague, 1 x Undisclosed location
  - [DE] 2 x Frankfurt
  - [SE] Stockholm
  - [UK] London
  - [IT] Milan
- **North America**
  - [US] 1 x California, 1 x Virginia
- **South America**
  - [BR] Sao Paulo
  - [CL] Santiago de Chile

17 locations

10 countries

4 continents

Upgrade results – number of servers

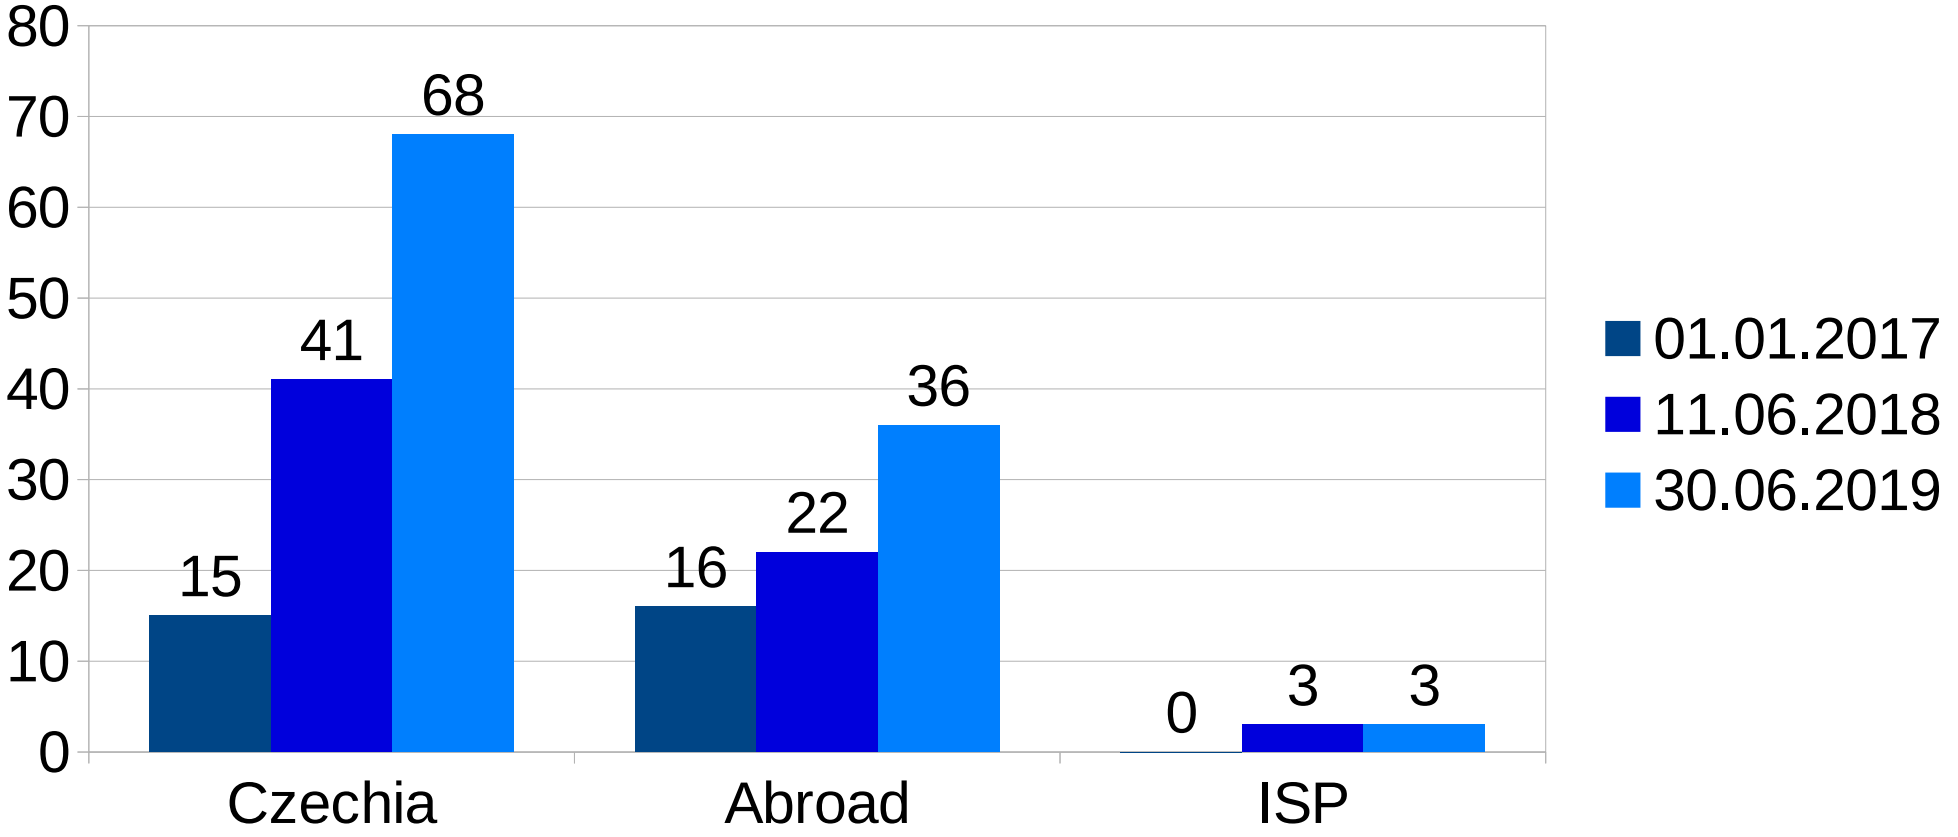

Upgrade results – network bandwitdth

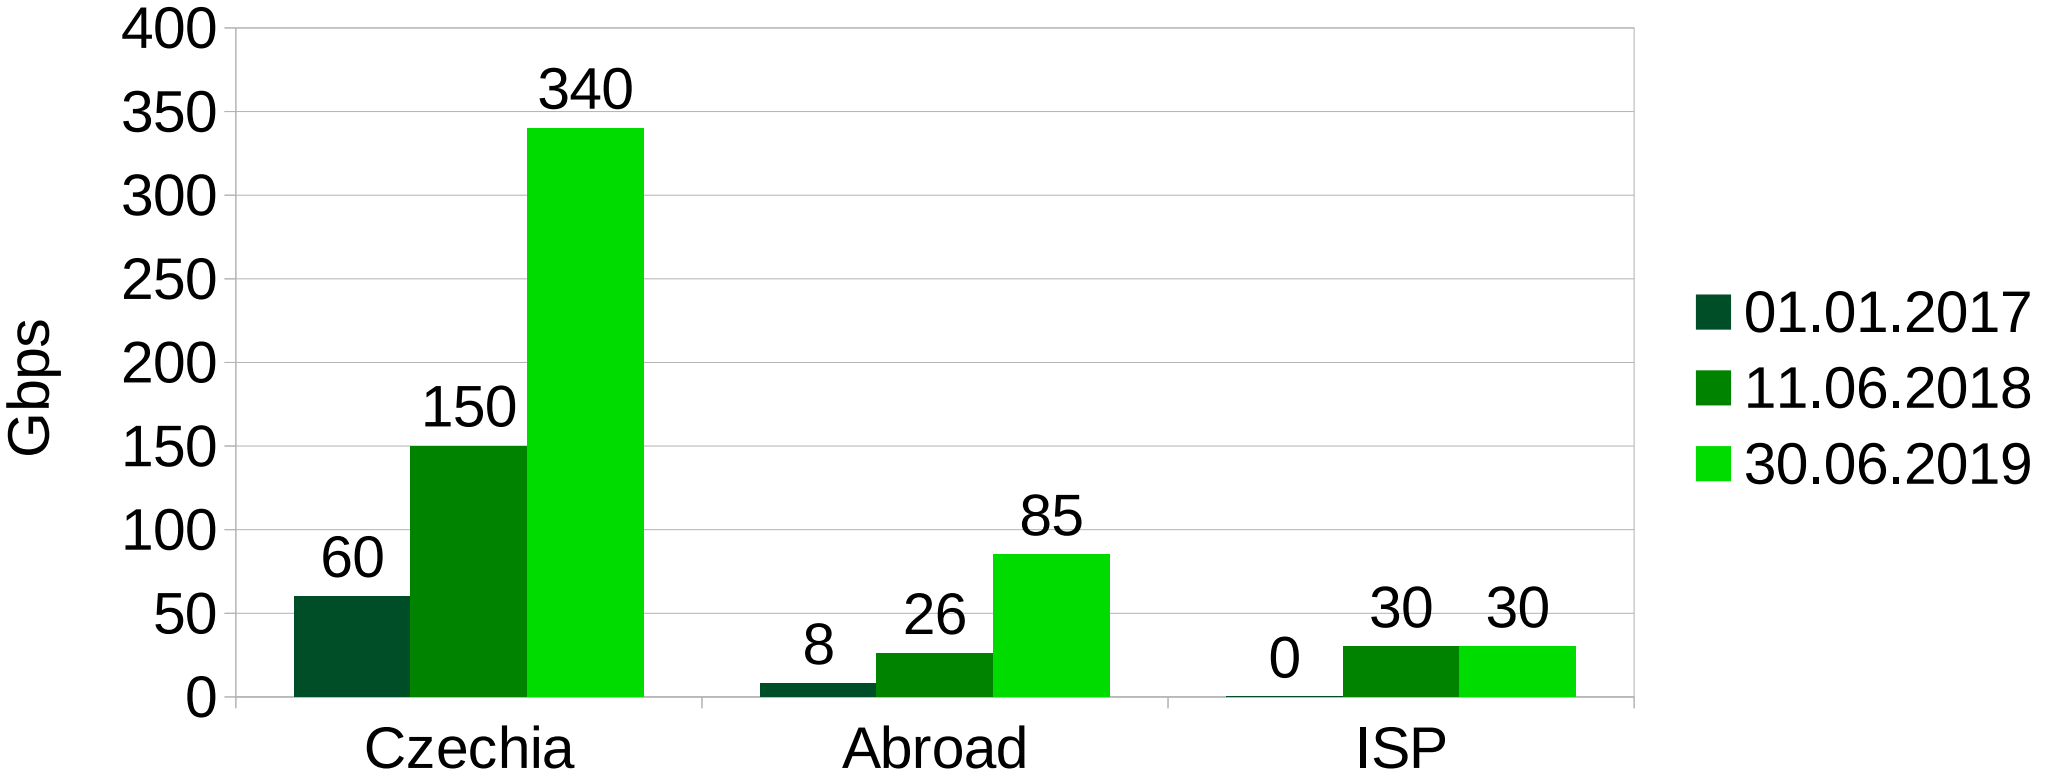

Upgrade results

- performance limits
  - 31 servers → **107 servers**
  - 20 000 000 → **200 000 000 QPS**
  - 60 Gbps → **400+ Gbps**

Not just performance - diversity

- HW vendors – Cisco, Juniper, Dell, HPE
- SW
  - OS (Ubuntu / Debian / OpenBSD)
  - BGP (**BIRD** / OpenBGPd / Quagga)
  - DNS (**KNOT** / BIND / NSD)
- Locations
  - Geographical
  - Different networks (IXP, **ISP**)

Upgrade results

1st 100 Gbps DNS stack

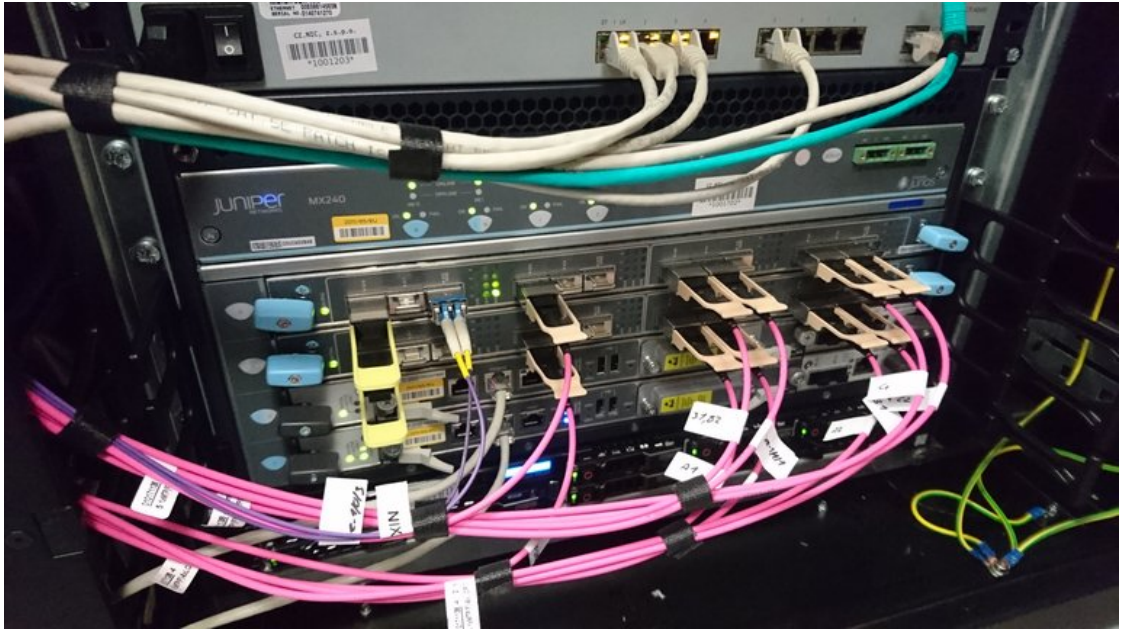

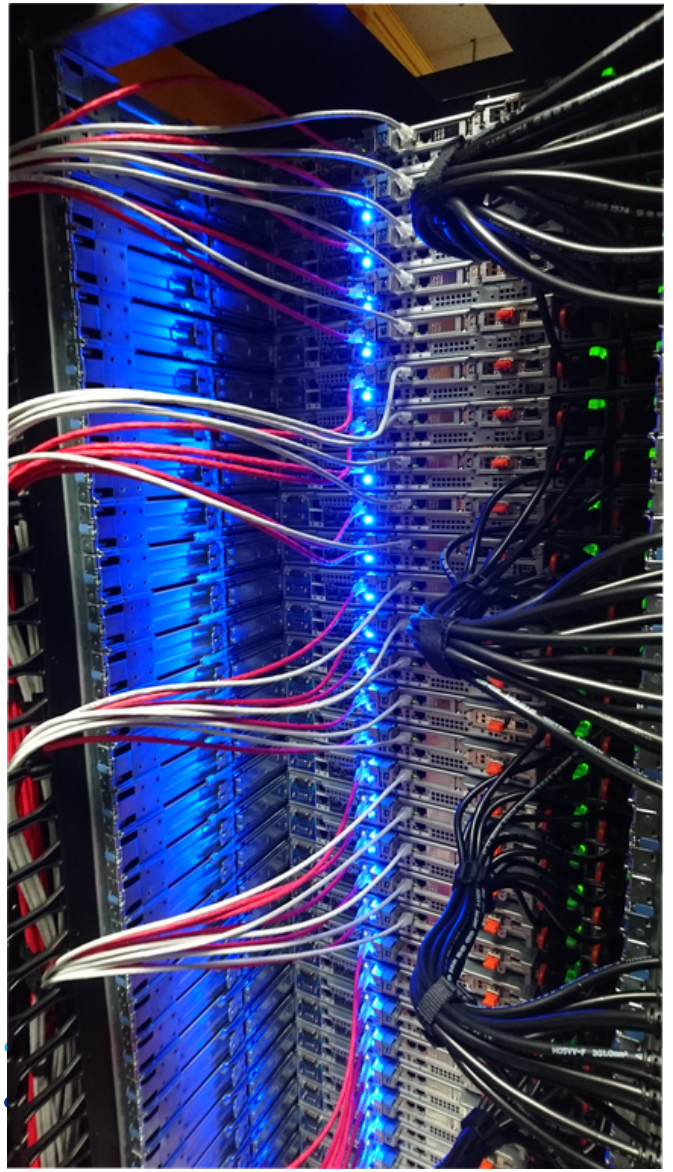

Upgrade results

2nd 100 Gbps DNS stack

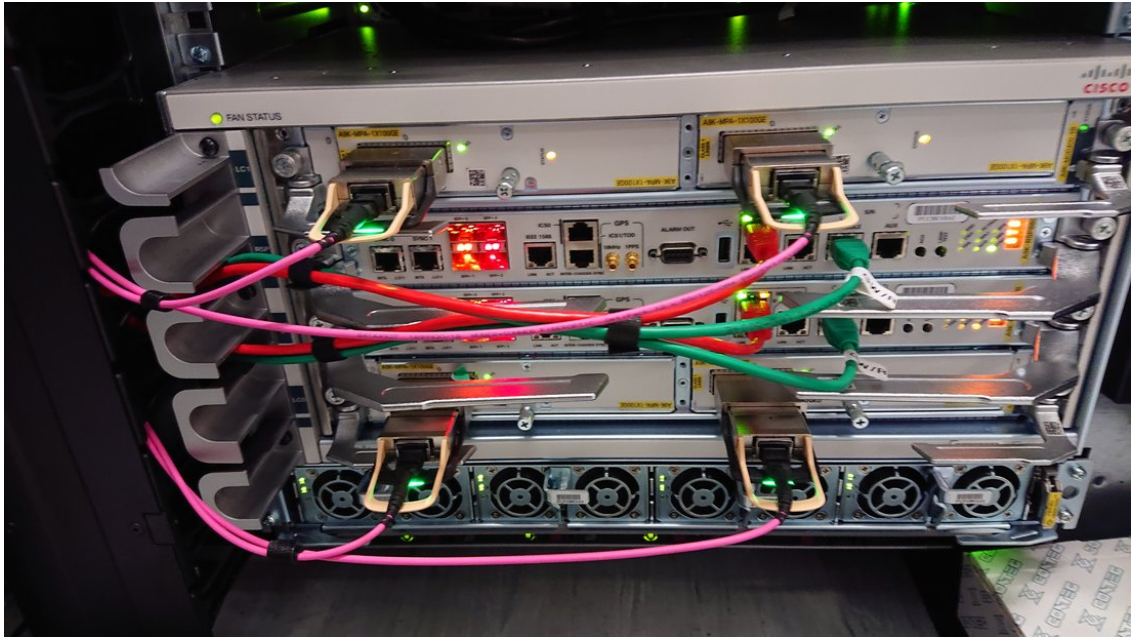

Sharing of the DNS infrastructure

Improving our anycast clouds

- Who sends queries to our servers?
- How long does it take for a query to reach our server?
  - typically active measurement (PING – DNS server vs client or probe)
  - we used ADAM (Advanced DNS Analysis and Monitoring) and passive analysis
    - captured DNS traffic from .CZ DNS anycast (1.-14.5.2019)
    - ~15,65 billion+ queries
      - UDP: 99.87% ~ 15,63 billion queries
      - **TCP: 0.13% ~ 21 million queries**

Improving our anycast clouds

- median RTT of a TCP handshake for each pair (client + server)
- evaluated RTT for each:
  - client
  - network
  - country

  ...

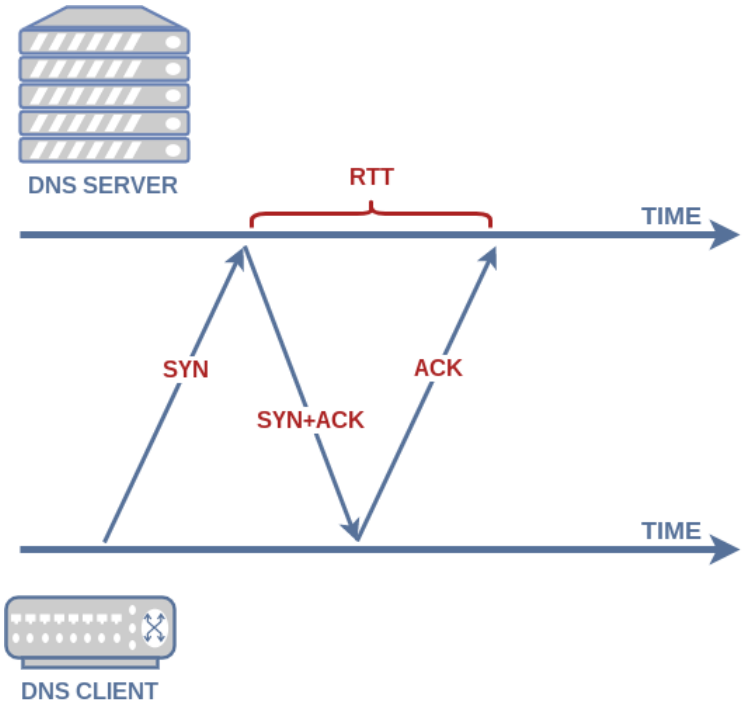

Improving our anycast clouds

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 – CE	**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

Improving our anycast clouds

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**	**0.0009**
217.31.193.164	CZ	25192	[Europe] CZ, Undisclosed	5171434	**12.7 ms**	**0.120**
217.31.193.164	CZ	25192	[Europe] CZ, Praha – CE	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**

Query distribution

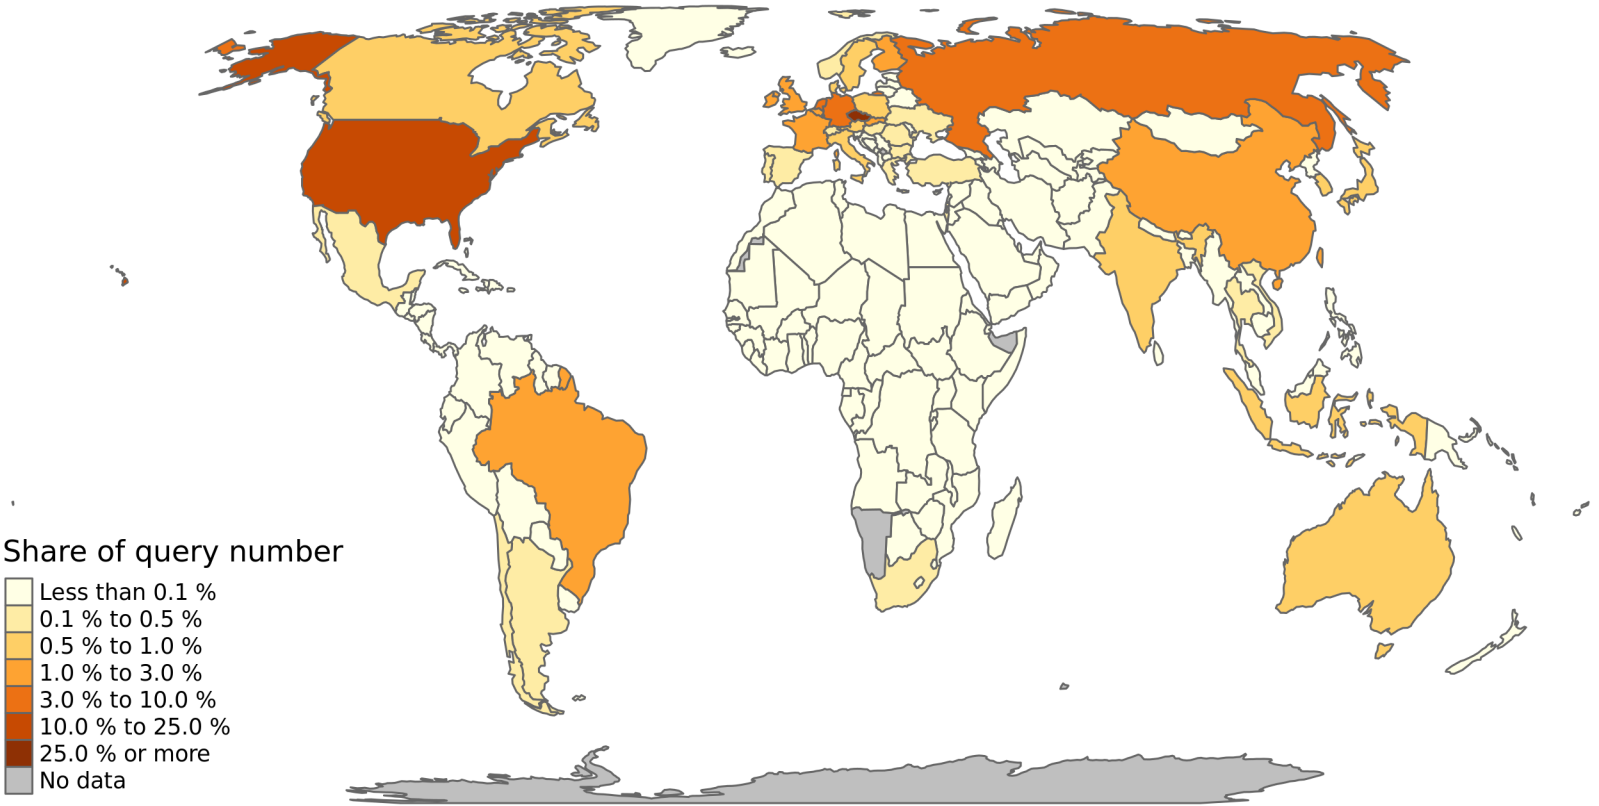

RTT by country

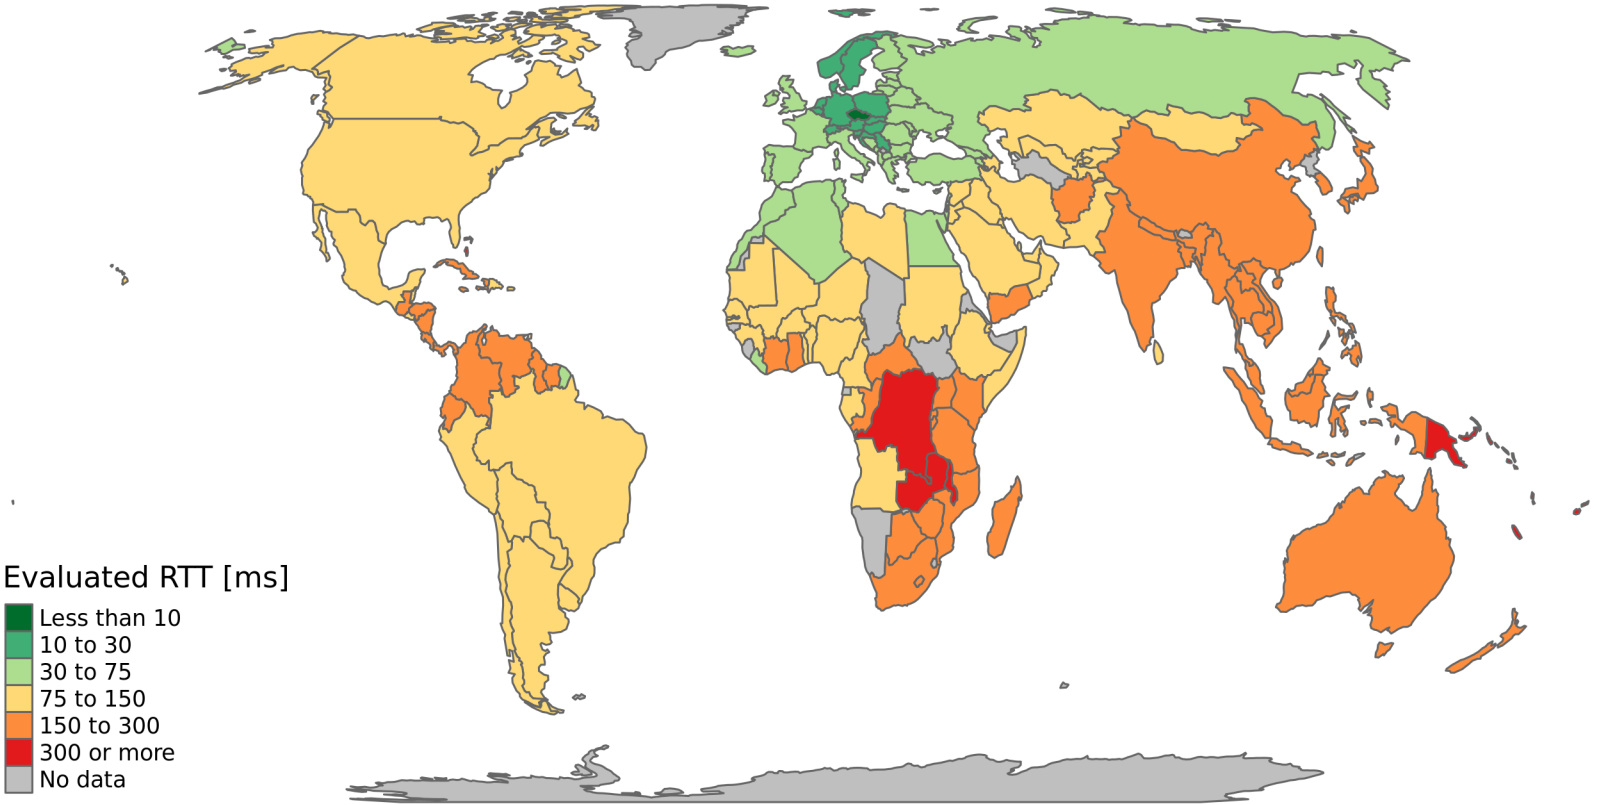

Queries vs RTT (Regions)

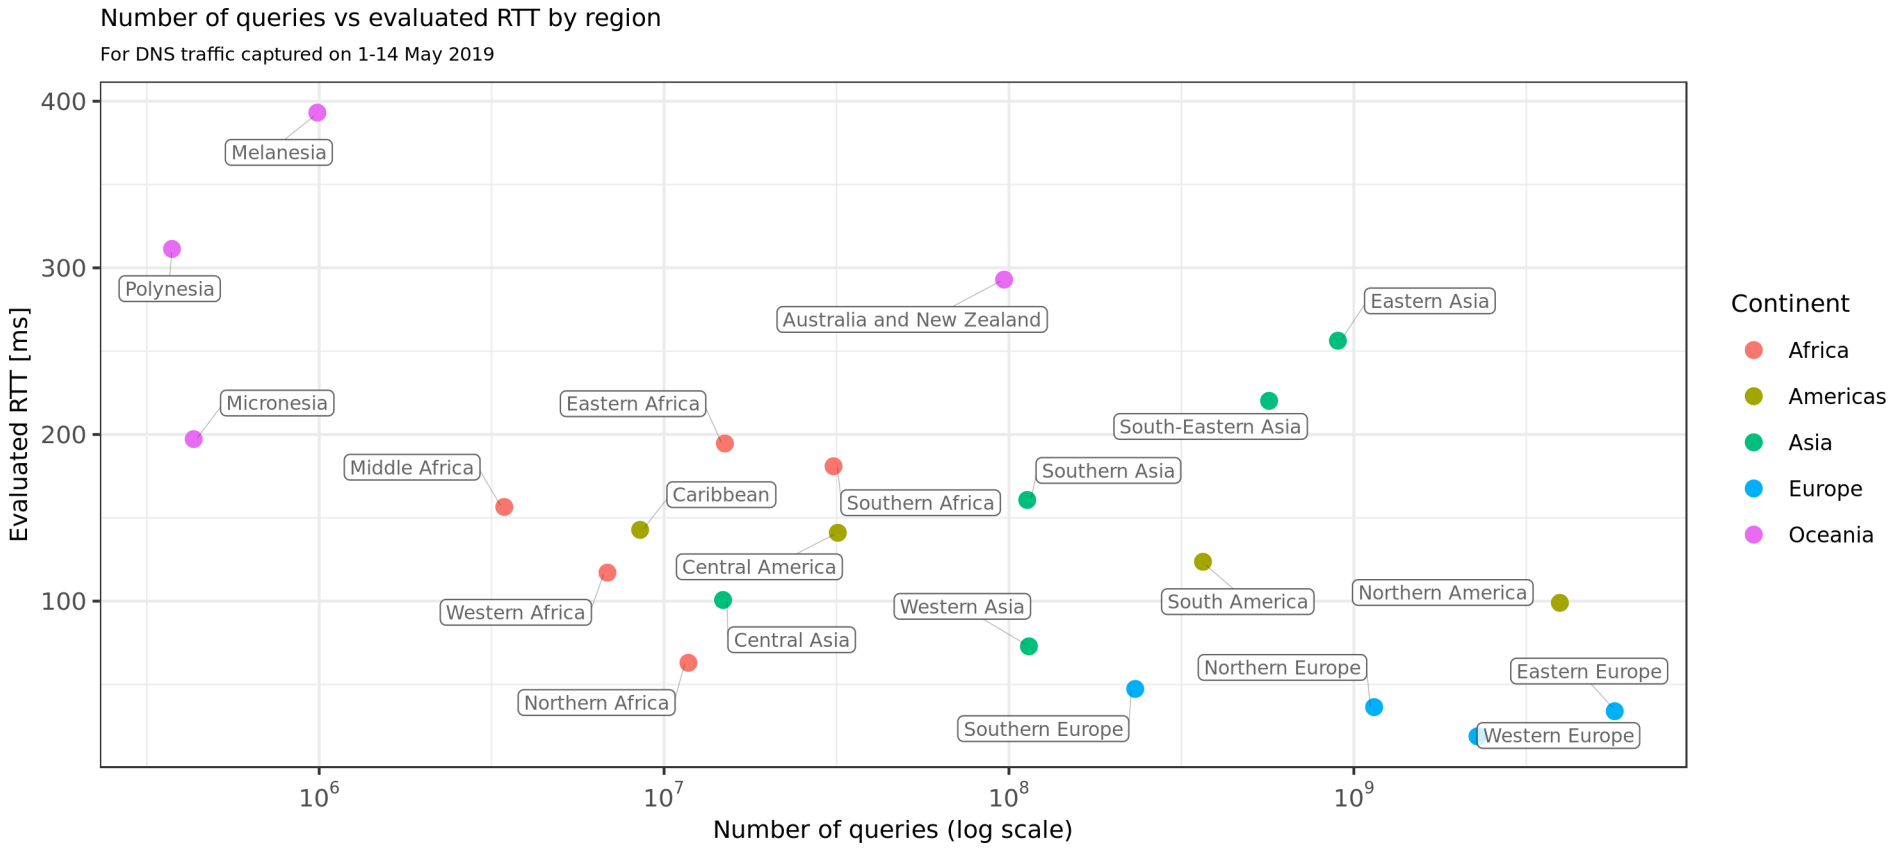

Queries vs RTT (TOP 50 Countries)

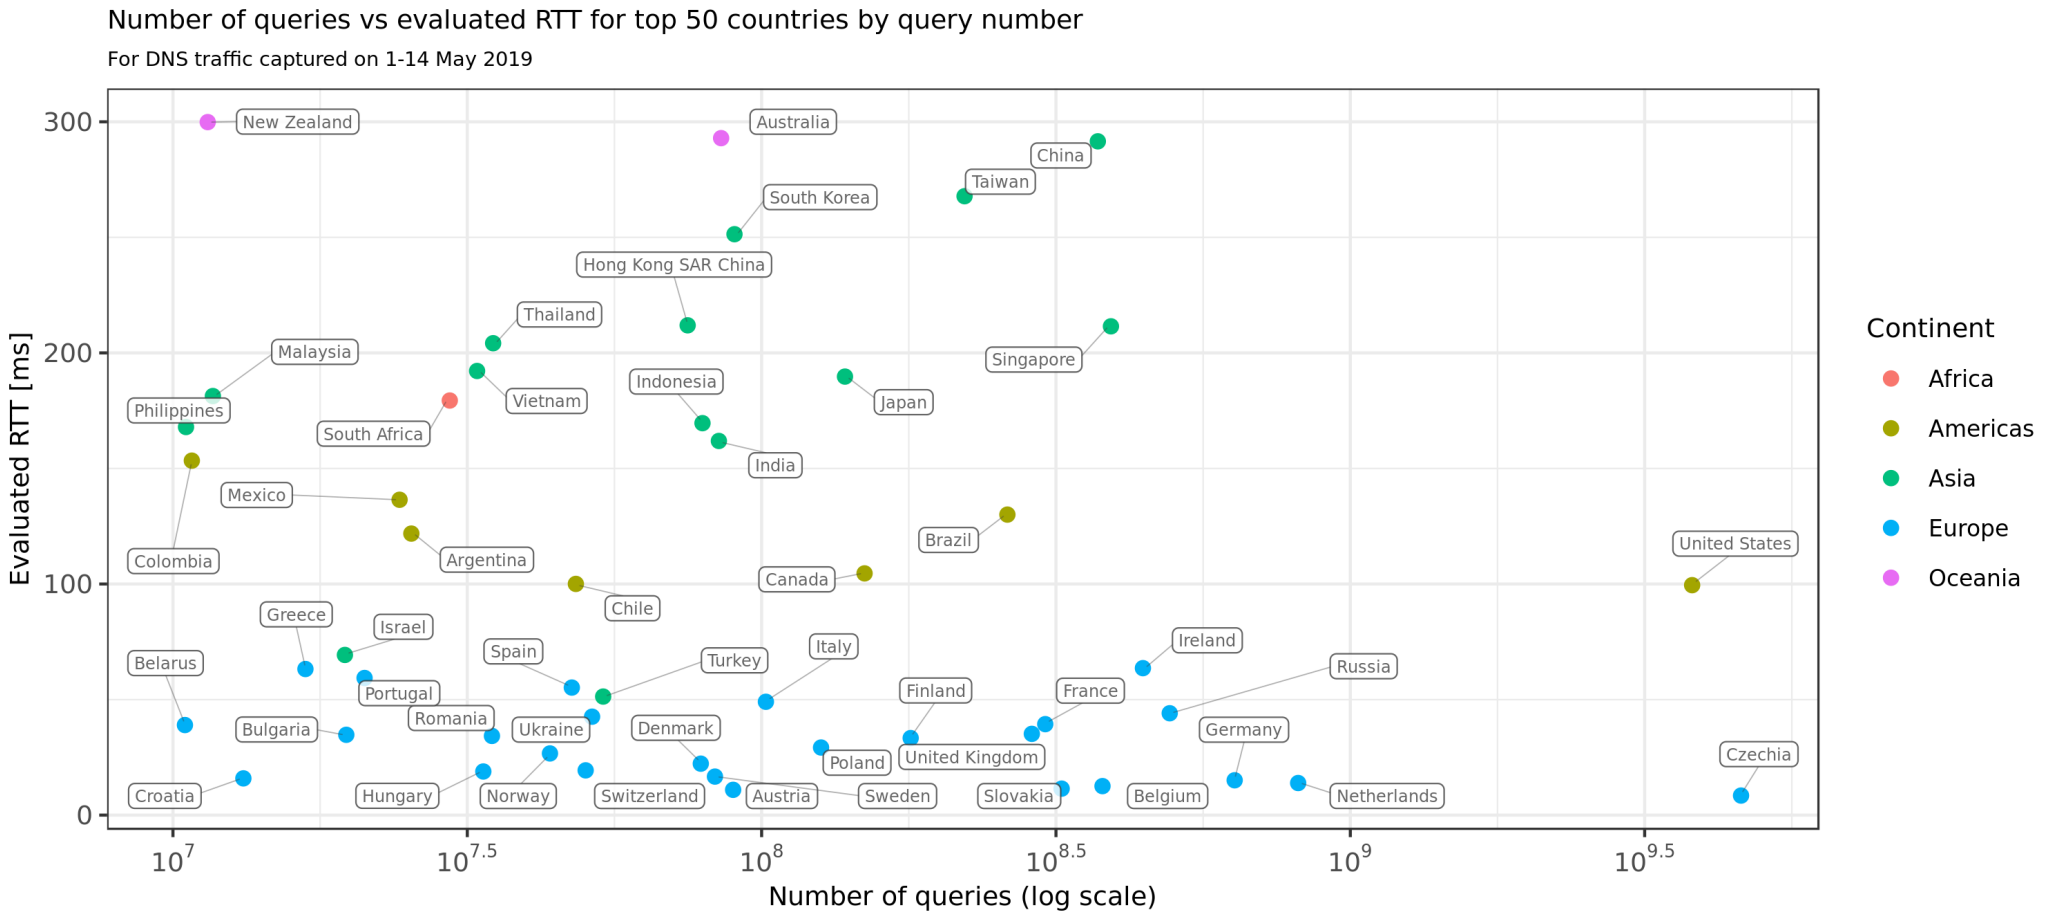

Traffic distribution (Example)

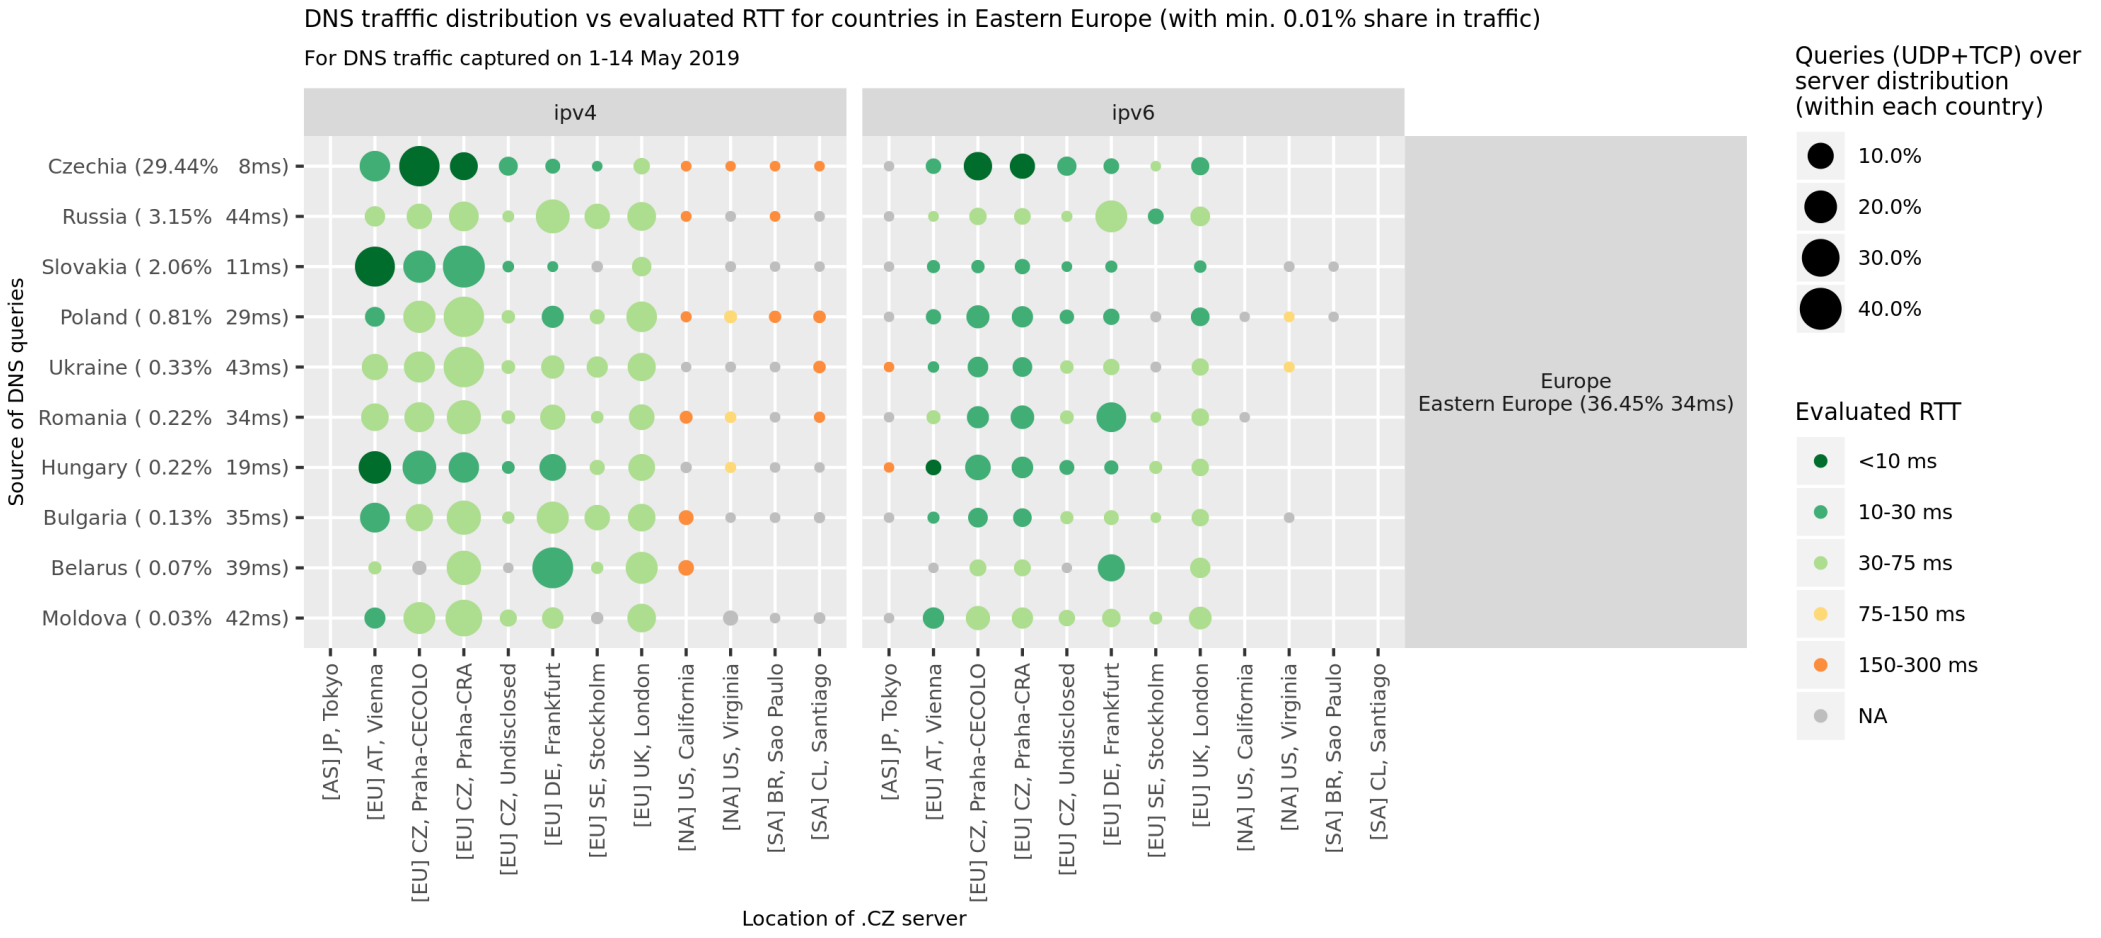

A DNS map from Czechia

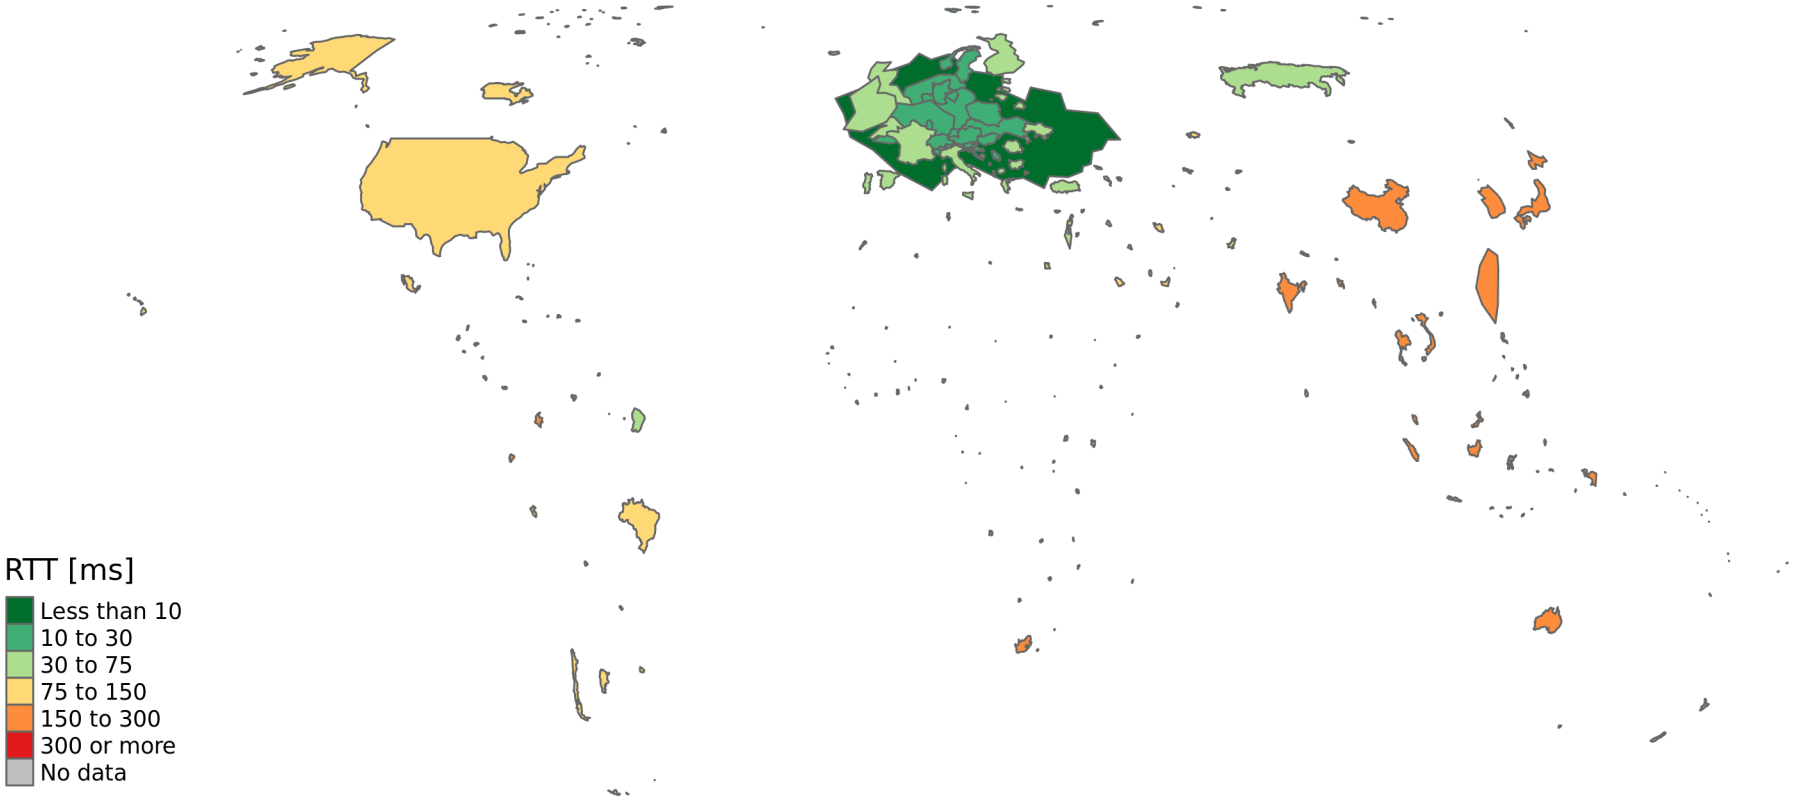

Conclusions

- Trying to prevent DDoS by brute force
- Monitoring using our tool set called ADAM
- Still expanding – looking for partners in NA, AP and AF
- And BTW evaluating creating own HW – combination of FPGA and many ARM mini servers

Thank you!

Ondřej Filip • ondrej.filip@nic.cz • https://www.nic.cz