In search of best locations for .CZ DNS servers

CSNOG 2023

Maciej Andziński • maciej.andzinski@nic.cz • 16.05.2023

CZ DOMAIN REGISTRY

Location of .CZ DNS servers

	City	Max QPS	Min QPS	Trend	Average
	Praha	6 110	4 063		5 143
	Frankfurt am Main	2 971	2 081		2 420
	Reston (VA)	2 434	1 608		1 943
	Seattle (WA)	2 061	913		1 778
	London	2 055	1 355		1 649
NEW ->	Singapore	1 648	999		1 168
	Tokyo	1 544	832		1 167
	undisclosed	1 243	592		802
	Wien	1 080	501		795
	Stockholm	1 231	614		766
	Milano	800	421		552
	São Paulo	745	375	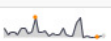	469
	Bratislava	476	44		378
	Santiago de Chile	659	110		340
NEW ->	Johannesburg	152	64		89

13 countries

5 continents

https://stats.nic.cz/dashboard/en/Traffic.html

CZ DOMAIN REGISTRY

What is the best location for our DNS servers?

- RTT-based approach
  - Goal: to improve the latency between a DNS client and .CZ DNS servers

What is the best location for our DNS servers?

- RTT-based approach
  - Identify sources of DNS queries
  - Measure the latency between a source and our sever

this is easy

this is difficult

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

- Use RTT of a TCP handshake to evaluate the latency between a DNS client and a DNS server

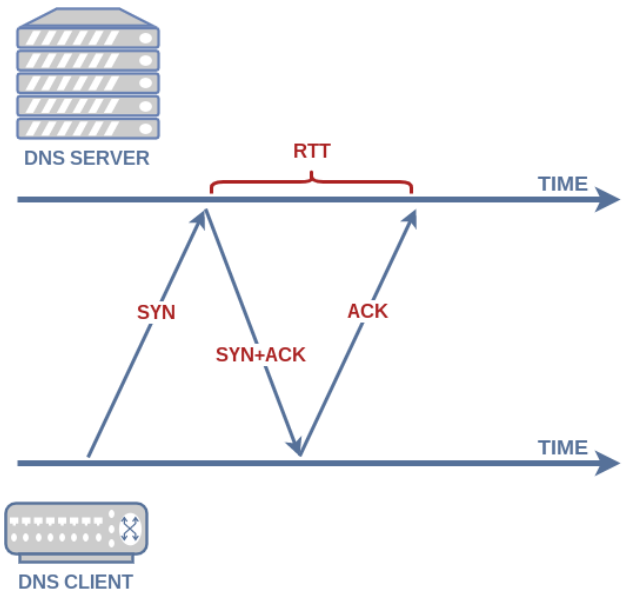

Our concept

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

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**

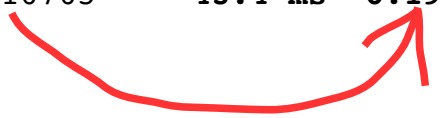

$$RTT = \sum_{i=1}^{n} Norm(w_i) \cdot RTT_i \qquad for\ RTT_i \neq NA$$

Evaluated RTT for 217.31.193.164 = **17.9 ms**

Evaluated RTT by country

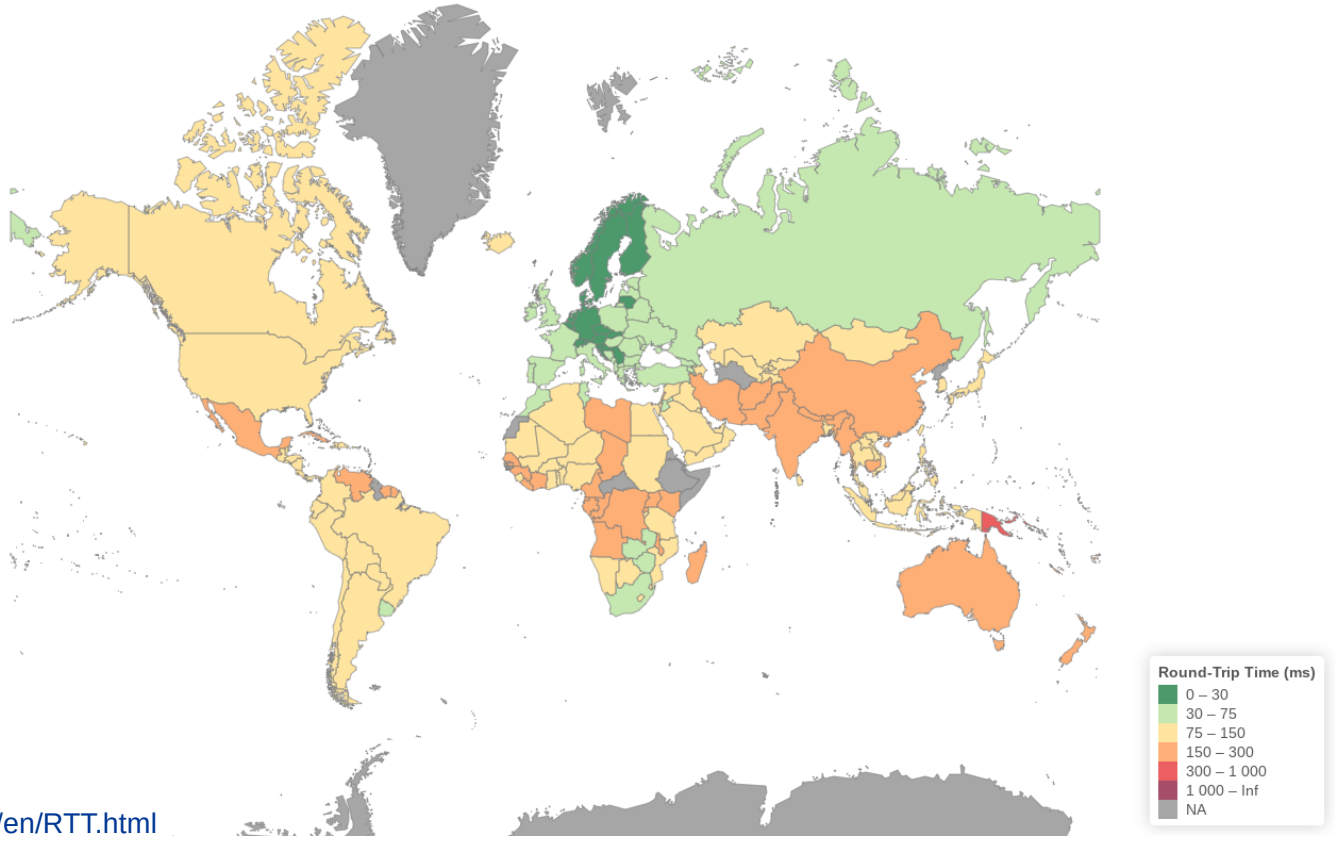

https://stats.nic.cz/dashboard/en/RTT.html

Previous RTT analyses in CZ.NIC

- Several studies since 2019
  - IT19 conference
    - https://www.youtube.com/watch?v=JDAxpqXHzY8
  - CENTR R&D workshop
    - https://centr.org/library/library/centr-event/rd14-andzinski-passive-analysis-of-dns-server-reachability-20190529.html
  - ADAM report: DNS RTT analysis reinforced
    - https://adam.pages.nic.cz/reports/adam/dns-tc-cz/
  - Master Thesis: Optimalizace provozu DNS anycastu pro .cz doménu (Lukáš Vacek)
    - https://dspace.cvut.cz/bitstream/handle/10467/87982/F8-DP-2020-Vacek-Lukas-thesis.pdf

Recent analysis

- A tool for selecting best location
  - Automatic
  - Easy to re-run
  - Producing easy-to-interpret results

A tool for selecting best location

1) Combine data:
   - RTT for ASNs
   - peeringDB data
2) Select best IX nodes suitable for installing new .CZ servers

A tool for selecting best location

- We introduce a measure called „RTT impact score“ (*RIS*)

$$RIS_{source} = QPS_{source} \cdot \left(RTT_{source} - RTT_{cz} \right)$$

  - *RIS>0* if a source makes the overall .CZ RTT worse
  - *RIS<0* if a source makes the overall .CZ RTT better

A tool for selecting best location

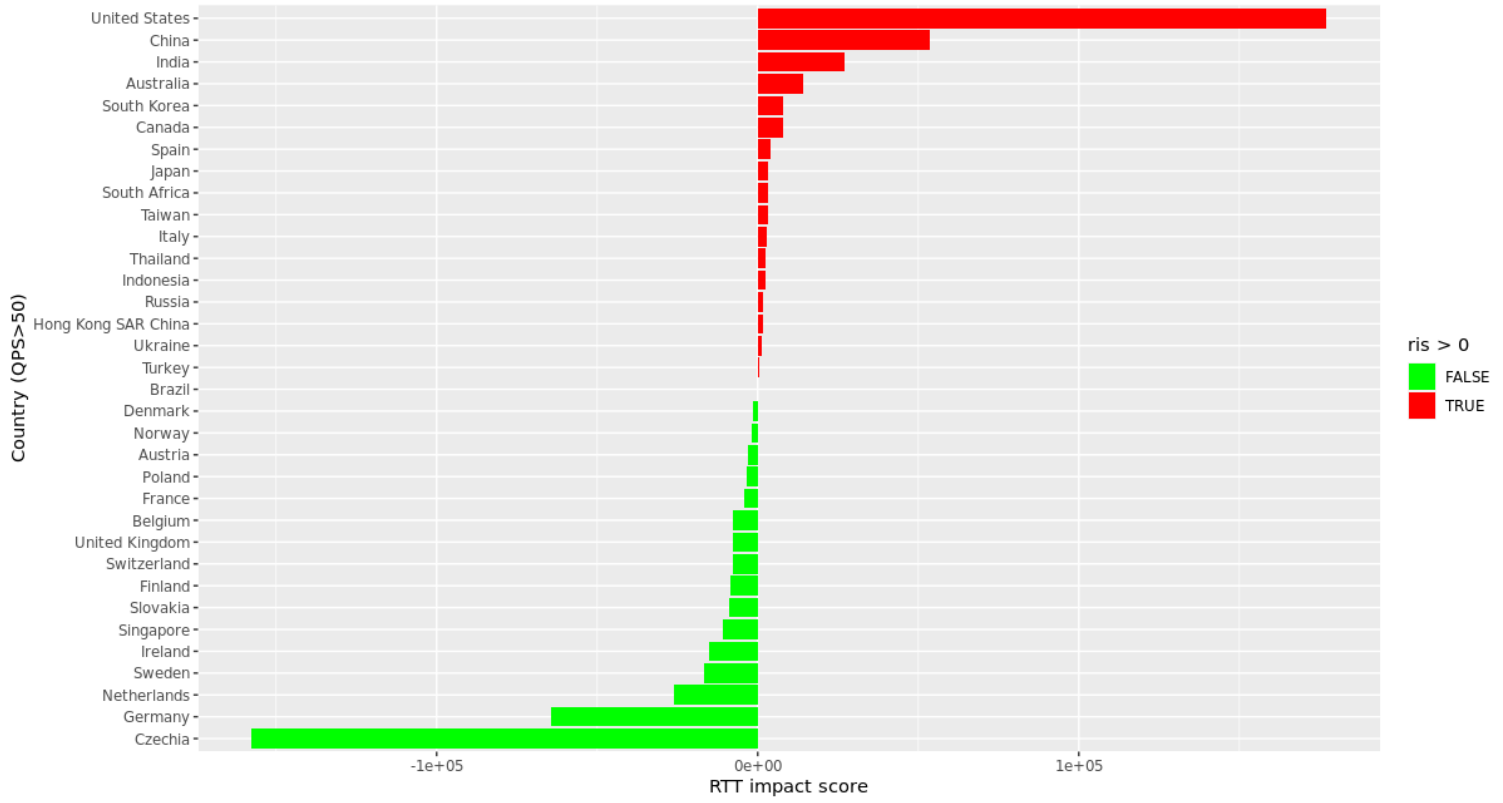

A tool for selecting best location

- „Voting“ method
  1) Select ASNs with *RIS>0*
  2) Sum *RIS* („votes“) for each IX node where selected ASNs peer
  3) Select IX nodes with highest number of votes
     1) Filter out IX nodes in countries with *RIS<0*
     2) Filter out IX nodes where *AS25192 CZ.NIC, z.s.p.o* already peers

Results

Suggested locations

List of suggested IXs for new .CZ servers location.

Show 25 entries

Search:

IX country	IX continent	IX region	IX description	RTT impact score	Covered ASNs in IX
United States	Americas	North America	NYIIX New York	344313	95
United States	Americas	North America	DE-CIX New York: DE-CIX New York Peering LAN	336472	100
India	Asia	South Asia	Extreme IX Mumbai: Extreme IX	309793	93
United States	Americas	North America	Equinix Ashburn	307878	143
United States	Americas	North America	Equinix Miami	302395	74
Hong Kong SAR China	Asia	East Asia & Pacific	Equinix Hong Kong	298989	98
United States	Americas	North America	Equinix Chicago	297682	112
Canada	Americas	North America	TorIX	293517	85
Japan	Asia	East Asia & Pacific	BBIX Tokyo	291554	132
India	Asia	South Asia	DE-CIX Mumbai: DE-CIX Mumbai Peering LAN	289241	129
United States	Americas	North America	Any2West	276930	137
Australia	Oceania	East Asia & Pacific	IX Australia (Sydney NSW): NSW-IX	274461	57
Japan	Asia	East Asia & Pacific	JPIX TOKYO	271090	109
Japan	Asia	East Asia & Pacific	JPNAP Tokyo: Peering	268956	102
Australia	Oceania	East Asia & Pacific	Equinix Sydney	266017	81

Sankey Diagram

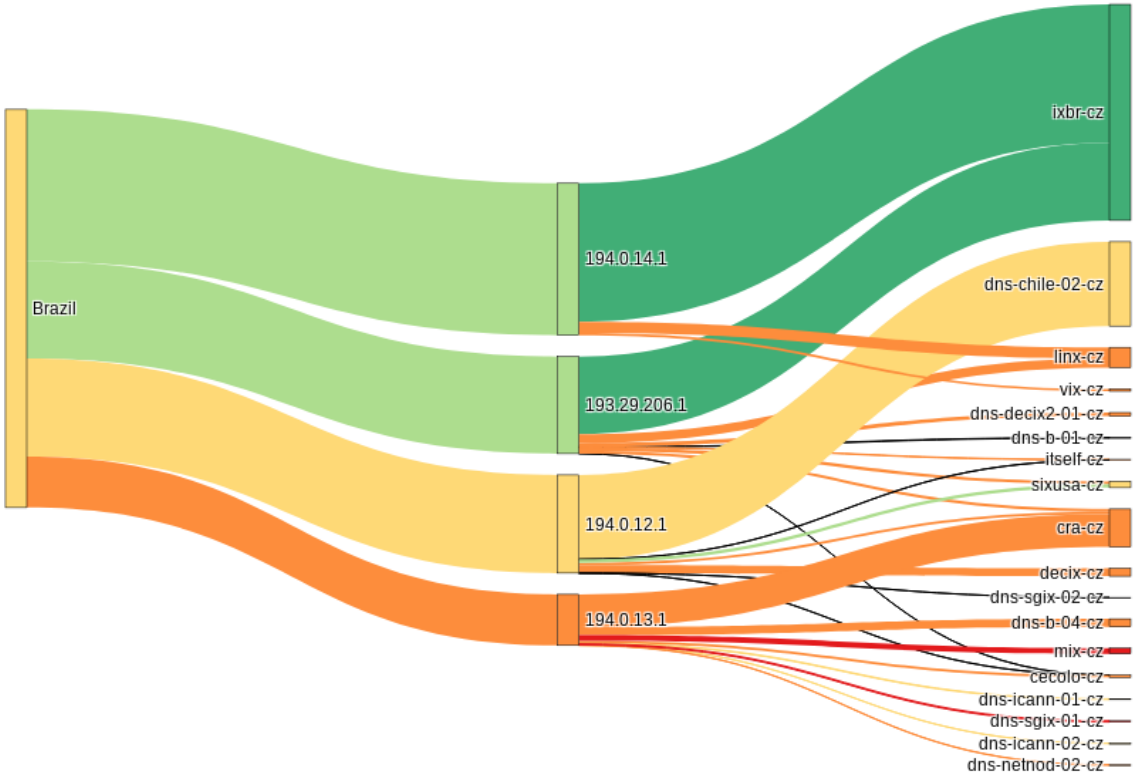

Conclusion

- RTT is just one of the aspects (but very important)
- „RTT impact score“ method
  - Helps to identify sources with negative impact on overall RTT
  - Not only RTT but also QPS matters
  - Can be powerful when combined with peeringDB data

Thank You

Maciej Andziński • maciej.andzinski@nic.cz

CZ DOMAIN REGISTRY