Clouding up the Internet: how centralized is DNS traffic becoming?

Giovane C. M. Moura (1), **Sebastian Castro (2)**, Wes Hardaker (3), Maarten Wullink (1) and Cristian Hesselman (1,4)

IRTF MAPRG - IETF 109

2020-11-16

Virtual Conference

1: SIDN Labs 2: InternetNZ 3: USC/ISI 4: University of Twente

Internet centralization concerns: US DOJ

The New York Times

Justice Department Opens Antitrust Review of Big Tech Companies

source: https://www.nytimes.com/2019/07/23/technology/justice-department-tech-antitrust.html

Centralization poses various risks

- Creates a **single point of failure**
- Privacy
- Market consolidation

DYN **DNS** 2016 Attack

source: https://www.nytimes.com/2016/10/22/business/internet-problems-attack.html

Can we measure Internet Centralization?

Easier said than done.

Measure it in terms of ?

- Users?
- Traffic?
- Networking infrastructure?
- Computing infrastructure?
- Market ?
- ...

Our approach:

- We focus on **DNS traffic**
- But **NOT** on *user* traffic
- We focus on traffic from resolvers to authoritative servers

Can we measure Internet Centralization?

Easier said than done.

Measure it in terms of ?

- Users?
- Traffic?
- Networking infrastructure?
- Computing infrastructure?
- Market ?
- ...

Our approach:

- We focus on **DNS traffic**
- But **NOT** on *user* traffic
- We focus on traffic from resolvers to authoritative servers

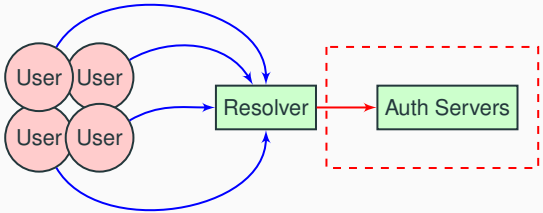

What we measure: DNS queries to

The Netherlands (.nl)

17.1M inhabitants
6M domain names (.nl)
Continent: Europe
Official language: Dutch

New Zealand (.nz)

4.8 M inhabitants
700k domain names (.nz)
Continent: Oceania
Official languages: English, Maori

B-Root

World
7.8 Billion inhabitants
1588 TLDs
Continents: 7
Language: *

What we measure: DNS queries from

From 5 Cloud/Content Providers

Company	ASes	Public DNS?
Google	15169	Yes
Amazon	7224, 8987, 9059, 14168, 16509	No
Microsoft	3598,6584, 8068–8075, 12076, 23468	No
Facebook	32934	No
Cloudflare	13335	Yes

So, what did we find?

Traffic to b.root-servers.net

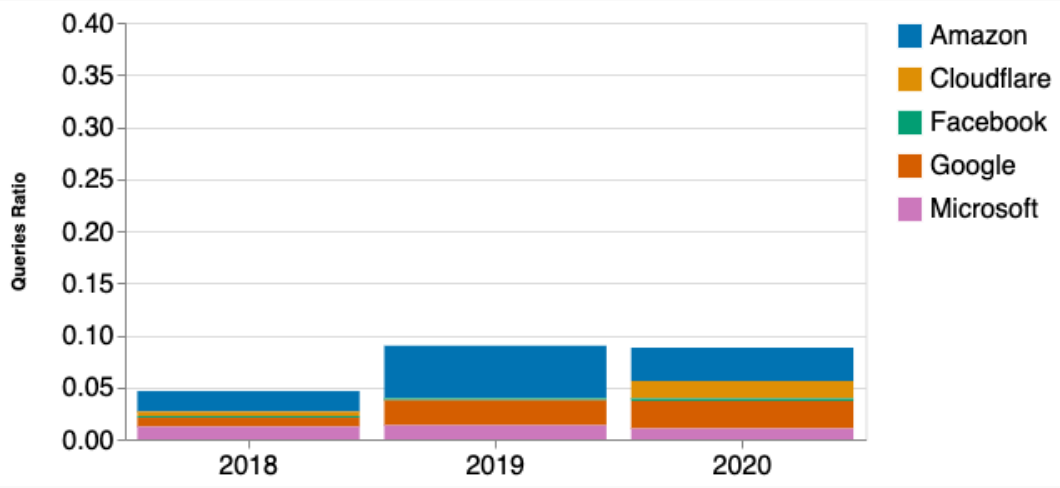

Traffic to .nz

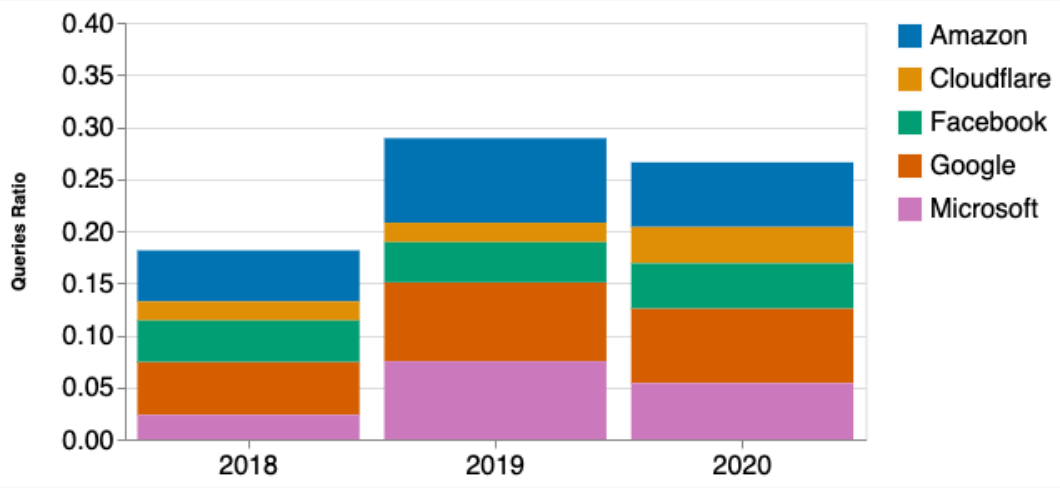

Traffic to .nl

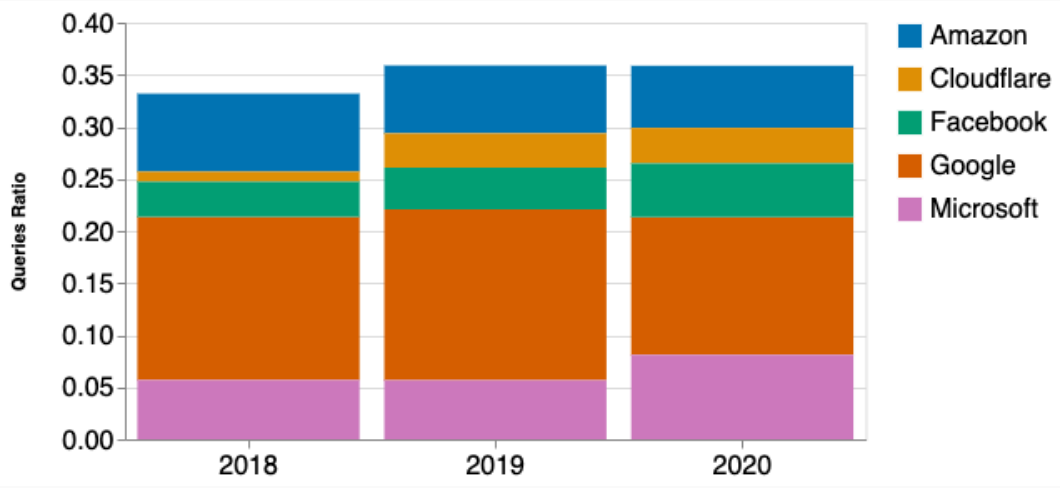

IPv4 vs IPv6 Adoption

- Roughly 50/50%:
  Google, Cloudflare
- More IPv6:
  Facebook (2019 onwards)
- Very little IPv6:
  Microsoft, Amazon

		.nl		.nz	
	Year	**IPv4**	**IPv6**	**IPv4**	**IPv6**
Google	2018	0.66	0.34	0.61	0.39
	2019	0.49	0.51	0.54	0.46
	2020	0.52	0.48	054	0.46
Amazon	2018	1	**0**	1	**0**
	2019	0.98	**0.02**	0.97	**0.03**
	2020	0.97	**0.03**	0.96	**0.04**
Microsoft	2018	1	**0**	1	**0**
	2019	1	**0**	1	**0**
	2020	1	**0**	1	**0**
Facebook	2018	0.52	0.48	0.51	0.49
	2019	0.24	0.76	0.19	0.81
	2020	0.24	0.76	0.17	0.83
Cloudflare	2018	0.54	0.46	0.54	0.46
	2019	0.57	0.43	0.56	0.44
	2020	0.51	0.49	0.49	0.51

IPv4 and IPv6 queries proportion

Traffic by query type to .nl

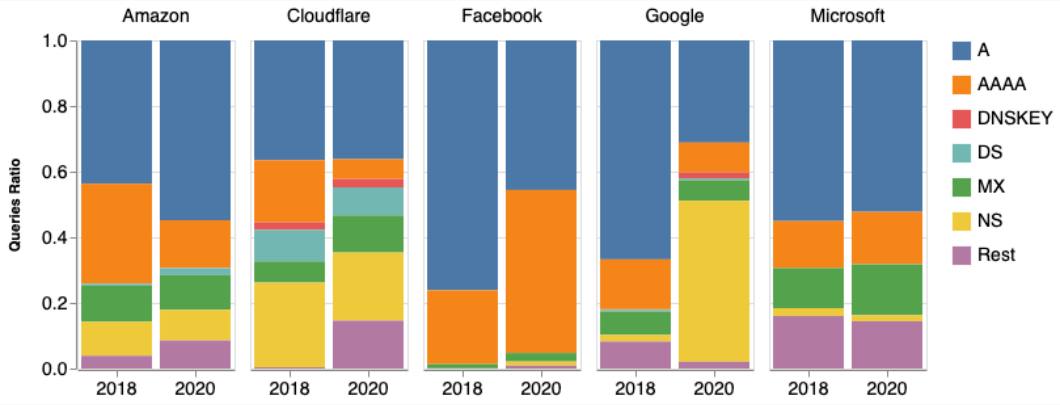

Our measurements revealed:

- Traffic levels differ per cloud provider
  - “Junk” queries vary by provider and year
- Query types sent vary significantly
  - By cloud provider
  - From year to year
- Key technology deployment variations
  - DNSSEC
  - IPv4 vs IPv6 usage
  - UDP vs TCP
  - Q-Name minimization
- Pros and Cons of centralization:
  - Rapid upgrades and rapid failures

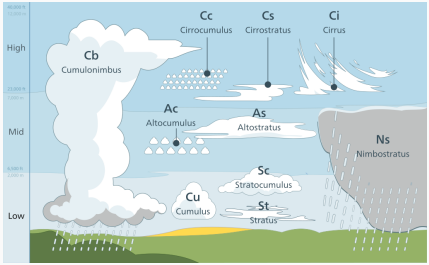

***Real-world* cloud types**

Paper (IMC2020):
Download it here
IMC 2020 Video Presentation:
Available here

Backup: Q-name minimization observations

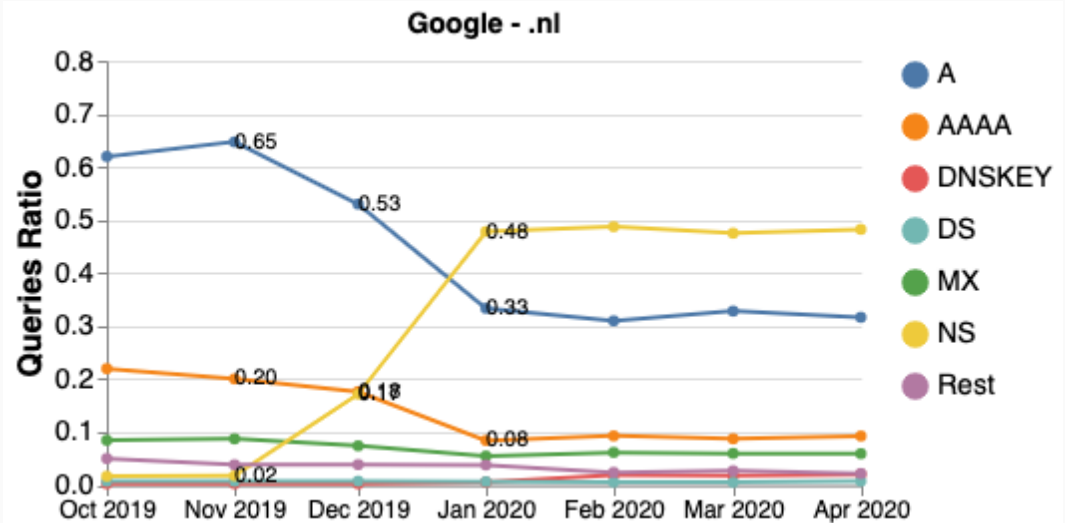

Backup: TCP vs UDP

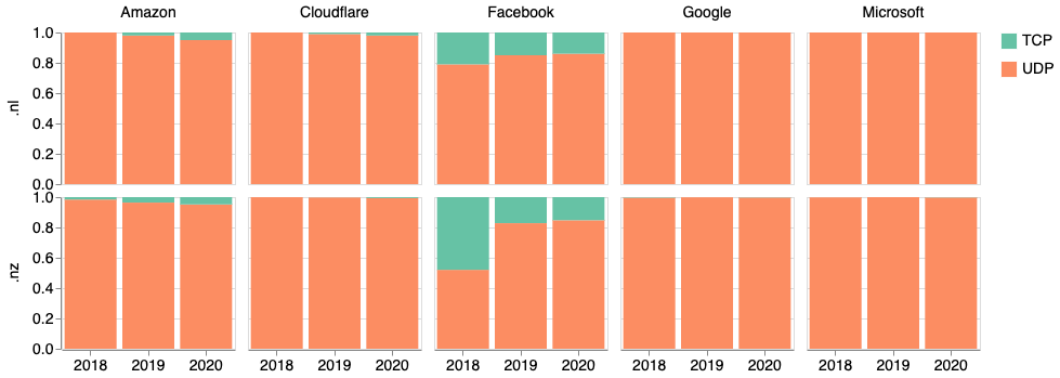

Backup: CDF on EDNS message size

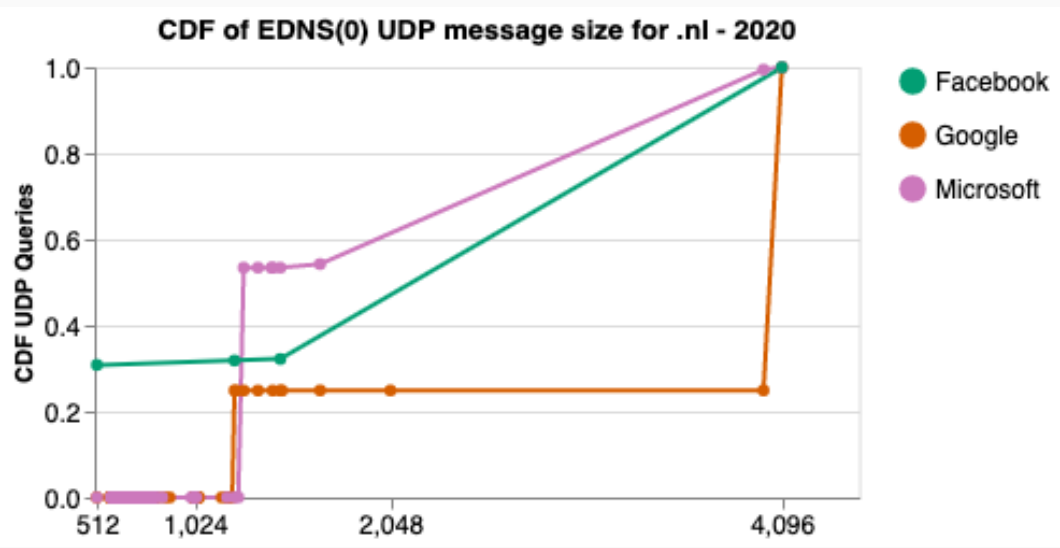

Backup: IPv4 vs IPv6