The Resolvers We Use

Joao Damas, Geoff Huston
APNIC

One of those wtf moments…

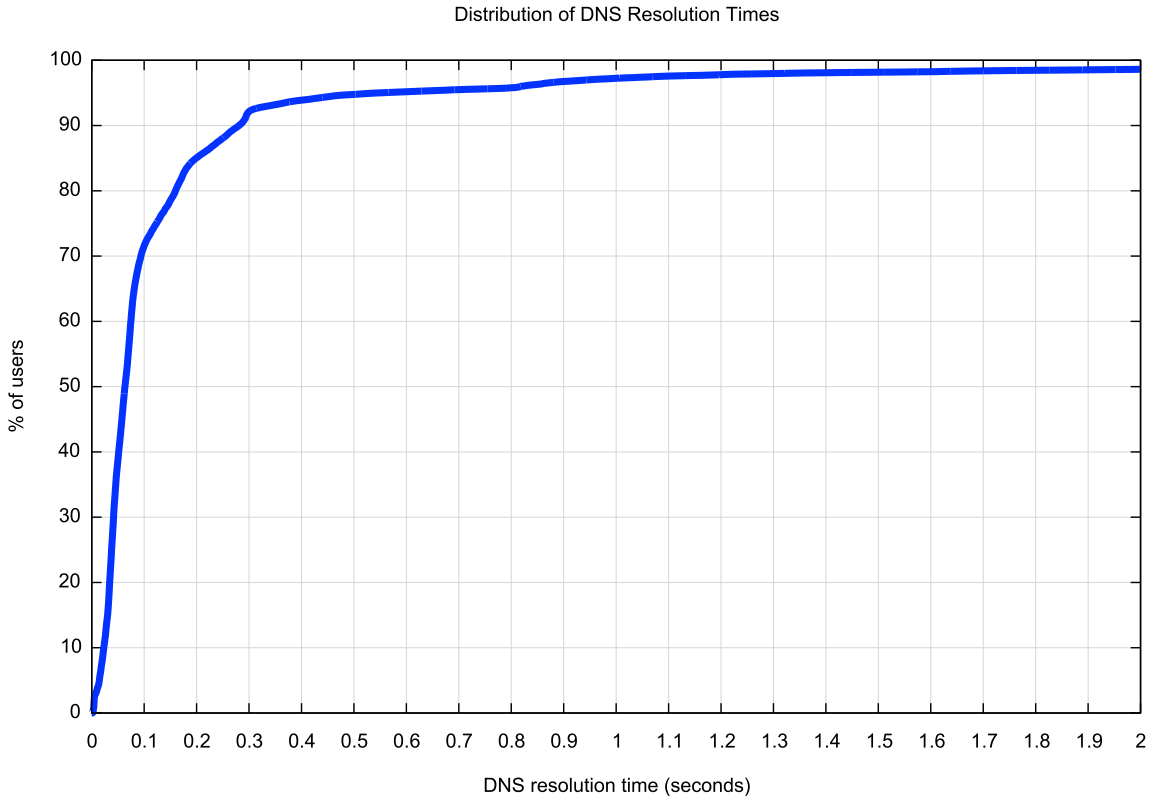

One of those wtf moments…

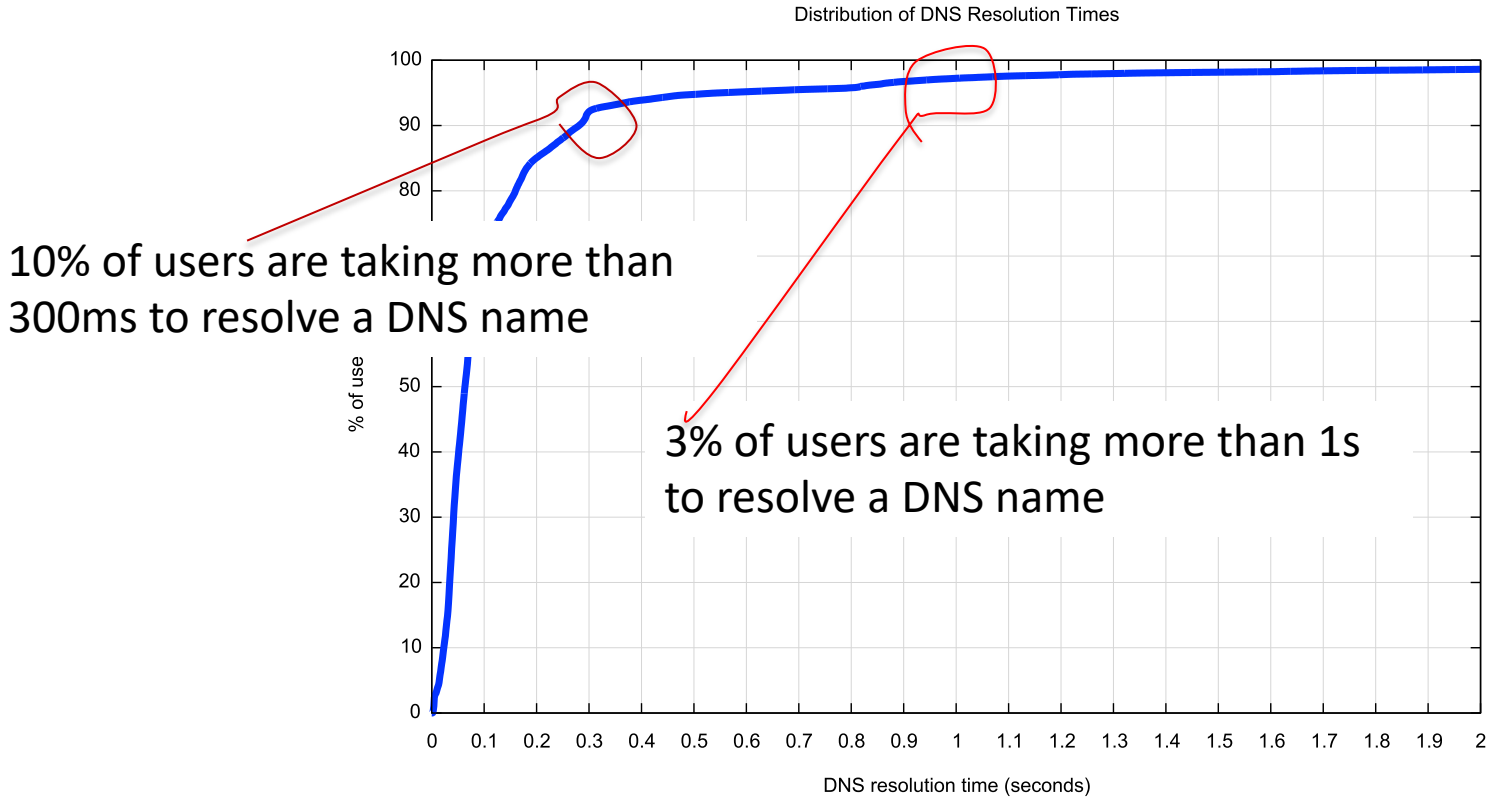

One of those wtf moments…

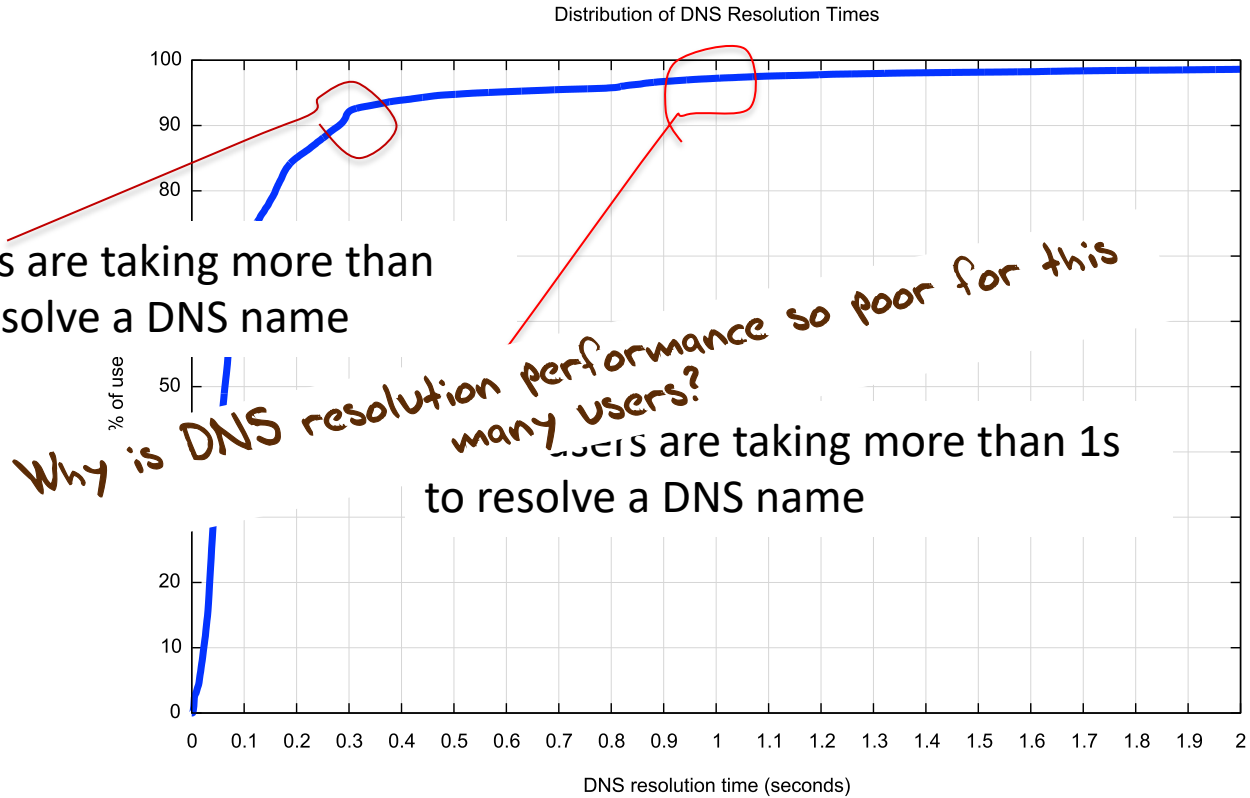

10% of users are taking more than 300ms to resolve a DNS name

Why is DNS resolution performance so poor for this many users?

are taking more than 1s to resolve a DNS name

And that leads to…

- It appears that what we think about how the DNS works and how folk actually are using the DNS is not well aligned
- We all think we understand how DNS resolution works in terms of the interchange of DNS protocol elements
  - But the performance of DNS name resolution depends on a number of additional factors, both in terms of the users' choice of resolvers and the name admin's choice of authoritative servers

But…

The issue is more than just a question of unexpected poor performance of name resolution.

There’s more to this...

Why is DNS resolution data valuable?

- Almost everything we do on the Internet starts with a DNS name resolution operation
- DNS resolver query logs contain a rich vein of real time information about what users are doing:

  query logs and can be analyzed to infer information about the users themselves through the names that their applications resolve

  query logs contains indirect pointers that can be used to derive aggregate aspects of users' demographics, preferences, purchases, etc

Information Leaks

The question of where your DNS query traffic is being sent is also a question of whether you are leaking a real time trail of your online activities

Which leads to an interesting question about today's Internet:

- To what extent is this DNS resolution data stream "leaked" outward?
  - Across network boundaries?
  - Across national boundaries?
- This second form of information leakage is "interesting"

  While many national regimes include regulations concerning personally identifying data, its not clear if these regulations extend these same protections to aliens who are not citizens of the country where the information is held

Measuring the Internet via its Users

At APNIC Labs we've been using online ads to measure the user's view of the Internet for some years

- We ask users to fetch a unique URL
- This involves a DNS resolution and a HTTP GET to our servers
- So we collect sets of DNS queries and WEB queries
- To see
  - how we are doing with the IPv6 transition
  - where DNSSEC validation is being used
  - And similar questions

Users and Resolvers

These data sets also allow us to match

- the IP address of the resolver that queries the authoritative name server (the “visible resolver”)

to

- the IP address of the client agent that retrieves the URL

Some Numbers

Ad Placements: 11M – 14M ad placements per day

Ad Distribution: Most countries have some 100K – 250K ad placements per day

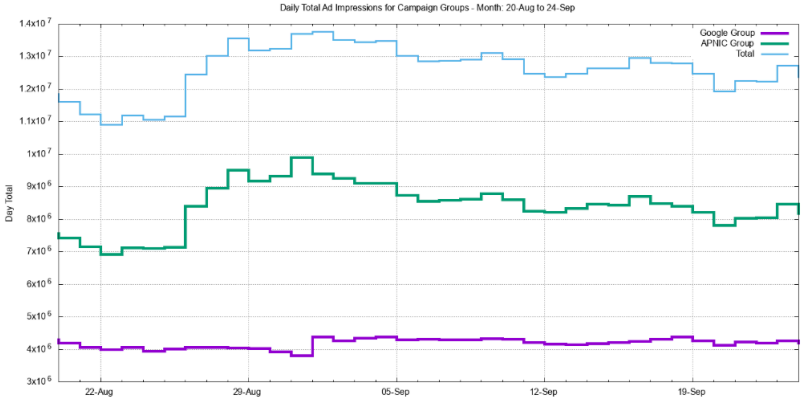

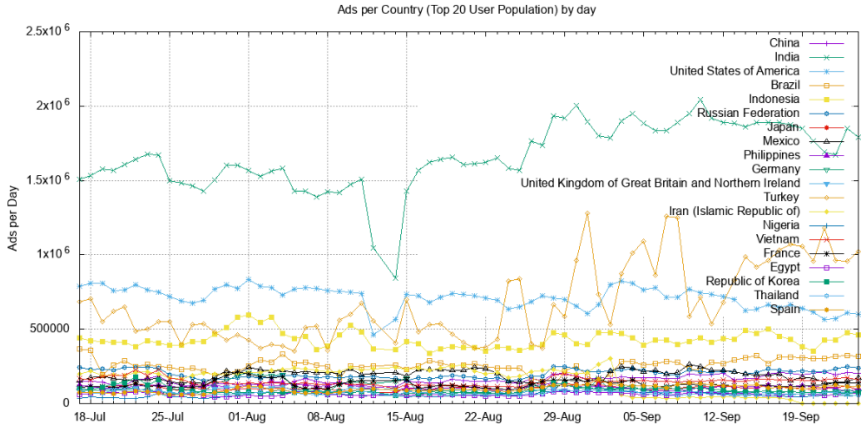

Top 25 Resolvers - By IP Address

Rank	Resolver	Use %	AS	AS Name
1	125.5.210.212	0.57%	AS7629	EPLDT, PH
2	196.188.52.8	0.49%	AS24757	EthioNet-AS, ET
3	202.56.215.67	0.34%	AS24560	Bharti Airtel, IN
4	2401:4900:50:9::5	0.34%	AS9498	Bharti Airtel, IN
5	129.205.112.254	0.28%	AS37148	Globa Com, NG
6	101.95.144.211	0.27%	AS4812	China Telecom, CN
7	2405:200:160c:1957:78::6	0.27%	AS55836	Reliance Jio, IN
8	49.45.29.22	0.27%	AS55836	Reliance Jio, IN
9	2405:200:160c:1957:78::4	0.27%	AS55836	Reliance Jio, IN
10	49.45.29.20	0.27%	AS55836	Reliance Jio, IN
11	49.45.29.21	0.27%	AS55836	Reliance Jio, IN
12	2405:200:160c:1957:78::5	0.26%	AS55836	Reliance Jio, IN
13	221.228.15.194	0.25%	AS4134	Chinanet Backbone, CN
14	101.95.144.210	0.25%	AS4812	China Telecom, CN
15	219.128.128.102	0.21%	AS58543	China Telecom, CN
16	2405:200:1613:1957:78::4	0.20%	AS55836	Reliance Jio, IN
17	49.44.189.220	0.20%	AS55836	Reliance Jio, IN
18	2405:200:1613:1957:78::5	0.20%	AS55836	Reliance Jio, IN
19	49.44.189.221	0.20%	AS55836	Reliance Jio, IN
20	49.44.189.222	0.20%	AS55836	Reliance Jio, IN
21	2405:200:1613:1957:78::6	0.20%	AS55836	Reliance Jio, IN
22	49.45.28.53	0.19%	AS55836	Reliance Jio, IN
23	2405:200:1609:1957:78::5	0.19%	AS55836	Reliance Jio, IN
24	2405:200:1609:1957:78::7	0.19%	AS55836	Reliance Jio, IN
25	49.45.28.55	0.19%	AS55836	Reliance Jio, IN

Top 25 Resolvers - By IP Address

Rank	Resolver	Use %	AS	AS Name
1	125.5.210.212	0.57%	AS7629	EPLDT, PH
2	196.188.52.8	0.49%	AS24757	EthioNet-AS, ET
3	202.56.215.67	0.34%	AS24560	Bharti Airtel, IN
4	2401:4900:50:9::5	0.34%	AS9498	Bharti Airtel, IN
5	129.205.112.254	0.28%	AS37148	Globa Com, NG
6	101.95.144.211	0.27%	AS4812	China Telecom, CN
7	2405:200:160c:1957:78::6	0.27%	AS55836	Reliance Jio, IN
8	49.45.29.22	0.27%	AS55836	Reliance Jio, IN
9	2405:200:160c:1957:78::4	0.27%	AS55836	Reliance Jio, IN
10	49.45.29.20	0.27%	AS55836	Reliance Jio, IN
11	49.45.29.21	0.27%	AS55836	Reliance Jio, IN
12	2405:200:160c:1957:78::5	0.26%	AS55836	Reliance Jio, IN
13	221.228.15.194	0.25%	AS4134	Chinanet Backbone, CN
14	101.95.144.210	0.25%	AS4812	China Telecom, CN
15	219.128.128.102	0.21%	AS58543	China Telecom, CN
16	2405:200:1613:1957:78::4	0.20%	AS55836	Reliance Jio, IN
17	49.44.189.220	0.20%	AS55836	Reliance Jio, IN
18	2405:200:1613:1957:78::5	0.20%	AS55836	Reliance Jio, IN
19	49.44.189.221	0.20%	AS55836	Reliance Jio, IN
20	49.44.189.222	0.20%	AS55836	Reliance Jio, IN
21	2405:200:1613:1957:78::6	0.20%	AS55836	Reliance Jio, IN
22	49.45.28.53	0.19%	AS55836	Reliance Jio, IN
23	2405:200:1609:1957:78::5	0.19%	AS55836	Reliance Jio, IN
24	2405:200:1609:1957:78::7	0.19%	AS55836	Reliance Jio, IN
25	49.45.28.55	0.19%	AS55836	Reliance Jio, IN

This list looks pretty strange!

A number of these resolvers share the same subnet – Are they different resolvers or part of a larger resolver "farm"?

Top Resolvers - by Origin AS

Rank	Resolver	Use %	Open Resolver / AS
1	Google DNS	9.39%	GOOGLE, US
2	AS55836	7.89%	Reliance Jio, IN
3	AS4134	5.22%	ChinaNET Backbone, CN
4	AS4837	2.86%	China Unicom, CN
5	AS9498	2.17%	Bharti Airtel, IN
6	AS9808	1.66%	China Mobile, CN
7	114dns	1.55%	ChinaNET Backbone, CN
8	OpenDNS	1.49%	OpenDNS, US
9	AS58543	1.47%	China Telecom, CN
10	AS24560	1.25%	Bharti Airtel, IN
11	dnspai	1.19%	China Telecom, CN
12	AS38266	1.10%	Vodafone India, IN
13	Onedns	1.01%	China Unicom Beijing Province Network, CN
14	AS8151	0.92%	Uninet, MX
15	AS45271	0.88%	Idea Cellular, IN
16	AS56040	0.83%	China Mobile, CN
17	AS7922	0.79%	Comcast, US
18	Cloudflare	0.76%	Cloudflare, US
19	Level3	0.76%	Level 3, US
20	AS23693	0.73%	Telekomunikasi Selular, ID
21	AS56046	0.71%	China MobilE, CN
22	AS9121	0.66%	TTNET, TR
23	AS17974	0.65%	Telekomunikasi Indonesia, ID
24	AS7629	0.63%	EPLDT, PH
25	AS132199	0.58%	Globe Telecom, PH

First Resolver or Full Resolver Set?

- End hosts are often configured with 2 or more recursive resolvers
- Is there much of a change in the use of recursive resolvers when we look at this full resolver set?

Top Resolvers - by Origin AS

Rank	Resolver	Use %	Open Resolver / AS
1	googlepdns	22.84%	Google, US
2	AS55836	7.92%	Reliance Jio Infocomm Limited, IN
3	AS4134	5.59%	ChinaNET Backbone, CN
4	AS4837	3.14%	China Unicom, CN
5	114dns	3.13%	ChinaNET, CN
6	AS9498	2.90%	Bharti Airtel IN
7	opendns	2.56%	OpenDNS, US
8	AS24560	2.54%	Bharti Airtel, IN
9	AS9808	1.87%	China Mobile, CN
10	level3	1.54%	Level 3, US
11	AS58543	1.53%	China Telecom, CN
12	dnspai	1.37%	China Telecom, CN
13	cloudflare	1.23%	Cloudflare, US
14	onedns	1.20%	China Unicom, CN
15	AS38266	1.10%	Vodafone India, IN
16	AS56046	1.05%	China Mobile, CN
17	AS8151	1.00%	Uninet, MX
18	AS56040	0.96%	China Mobile, CN
19	AS45271	0.91%	Idea Cellular, N
20	AS7922	0.80%	Comcast, US
21	AS23693	0.73%	Telekomunikasi Selular, ID
22	AS7629	0.70%	EPLDT, PH
23	AS9121	0.68%	TTNET, TR
24	AS17974	0.66%	Telekomunikasi Indonesia, ID
25	AS132199	0.58%	Globe Telecom, PH

Full Resolver Set

Resolver Distribution

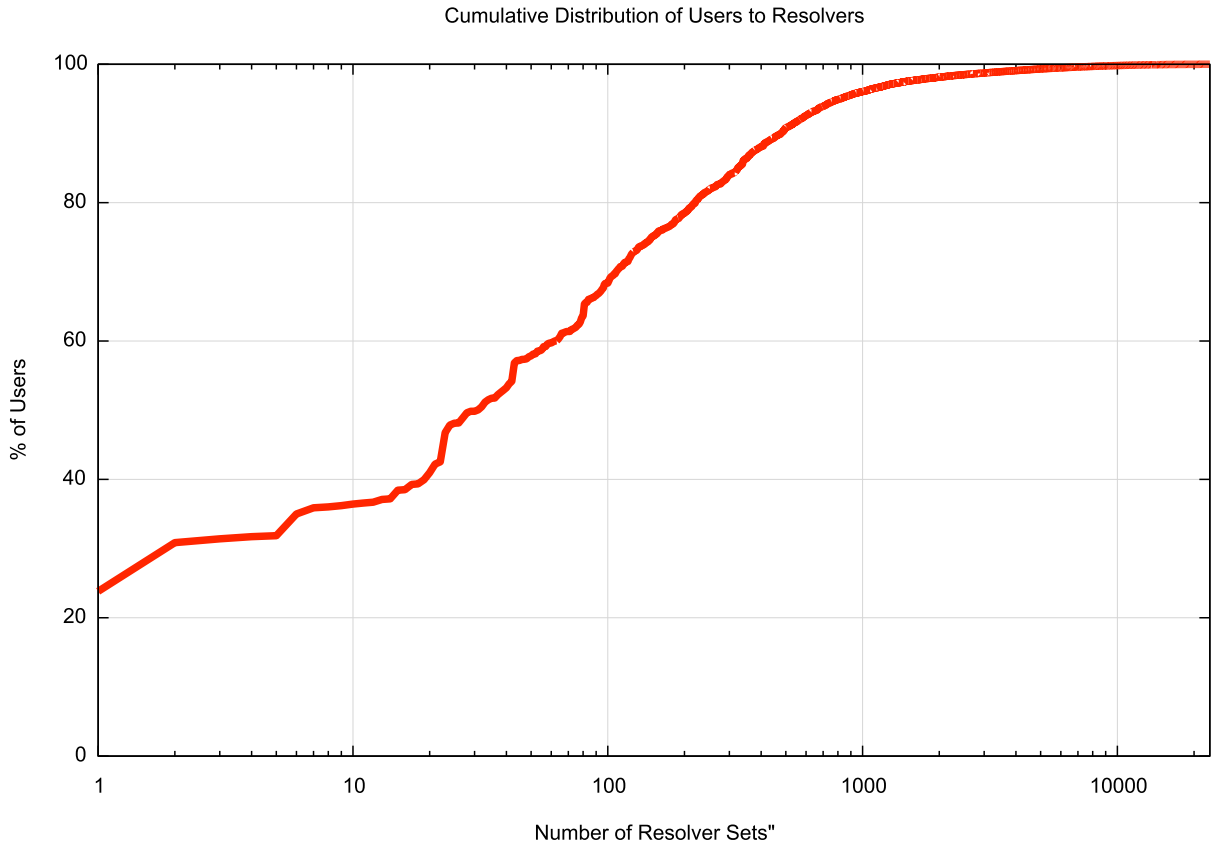

Resolver Distribution

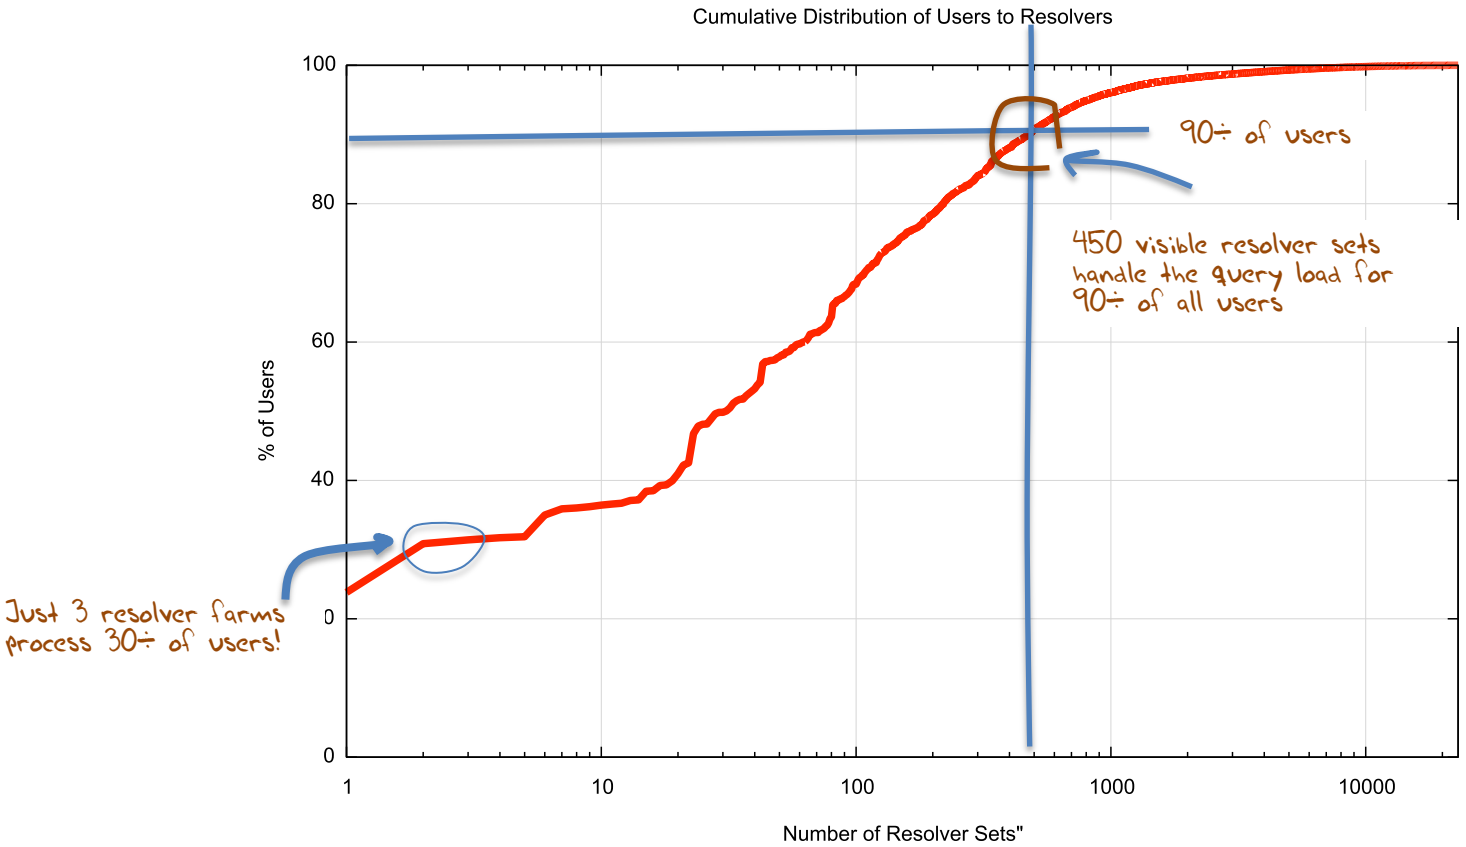

Counting Resolver Use

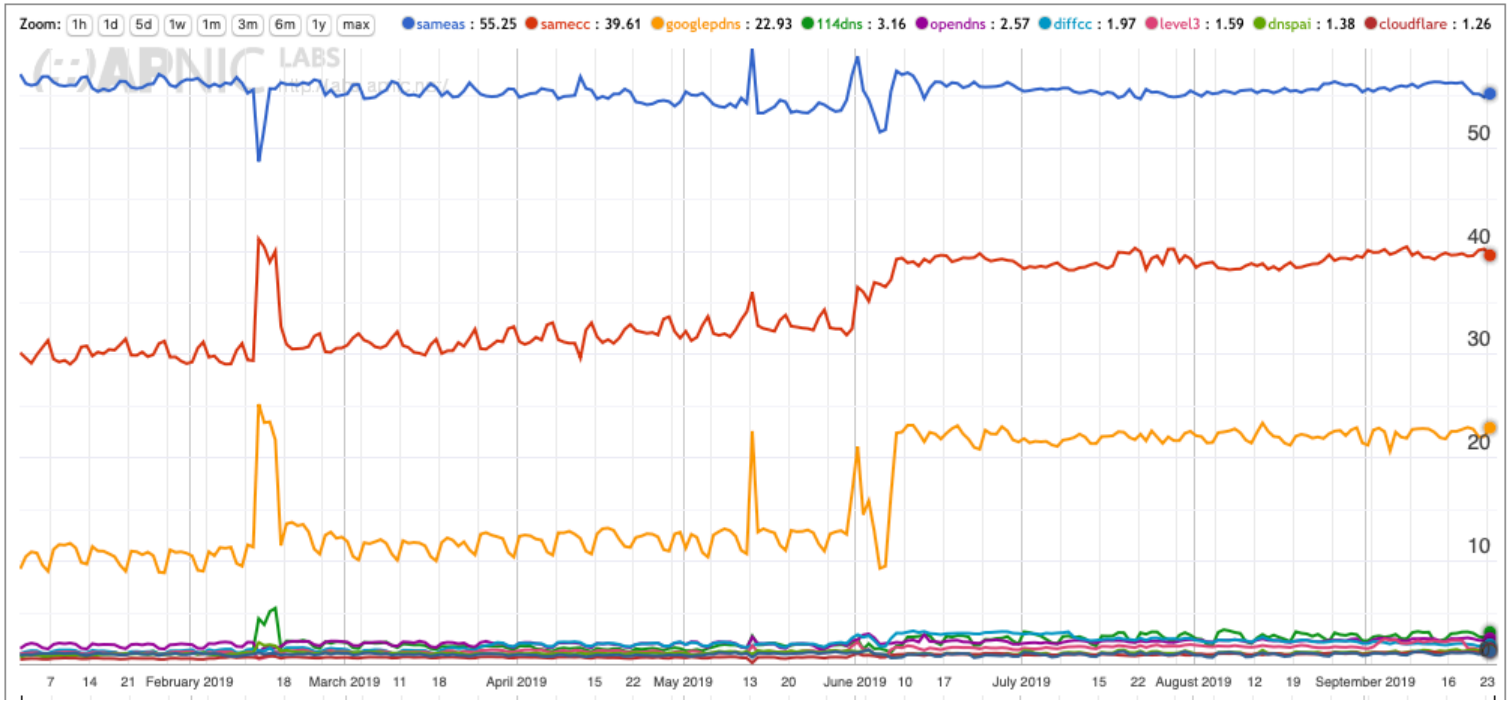

Counting Resolver Use

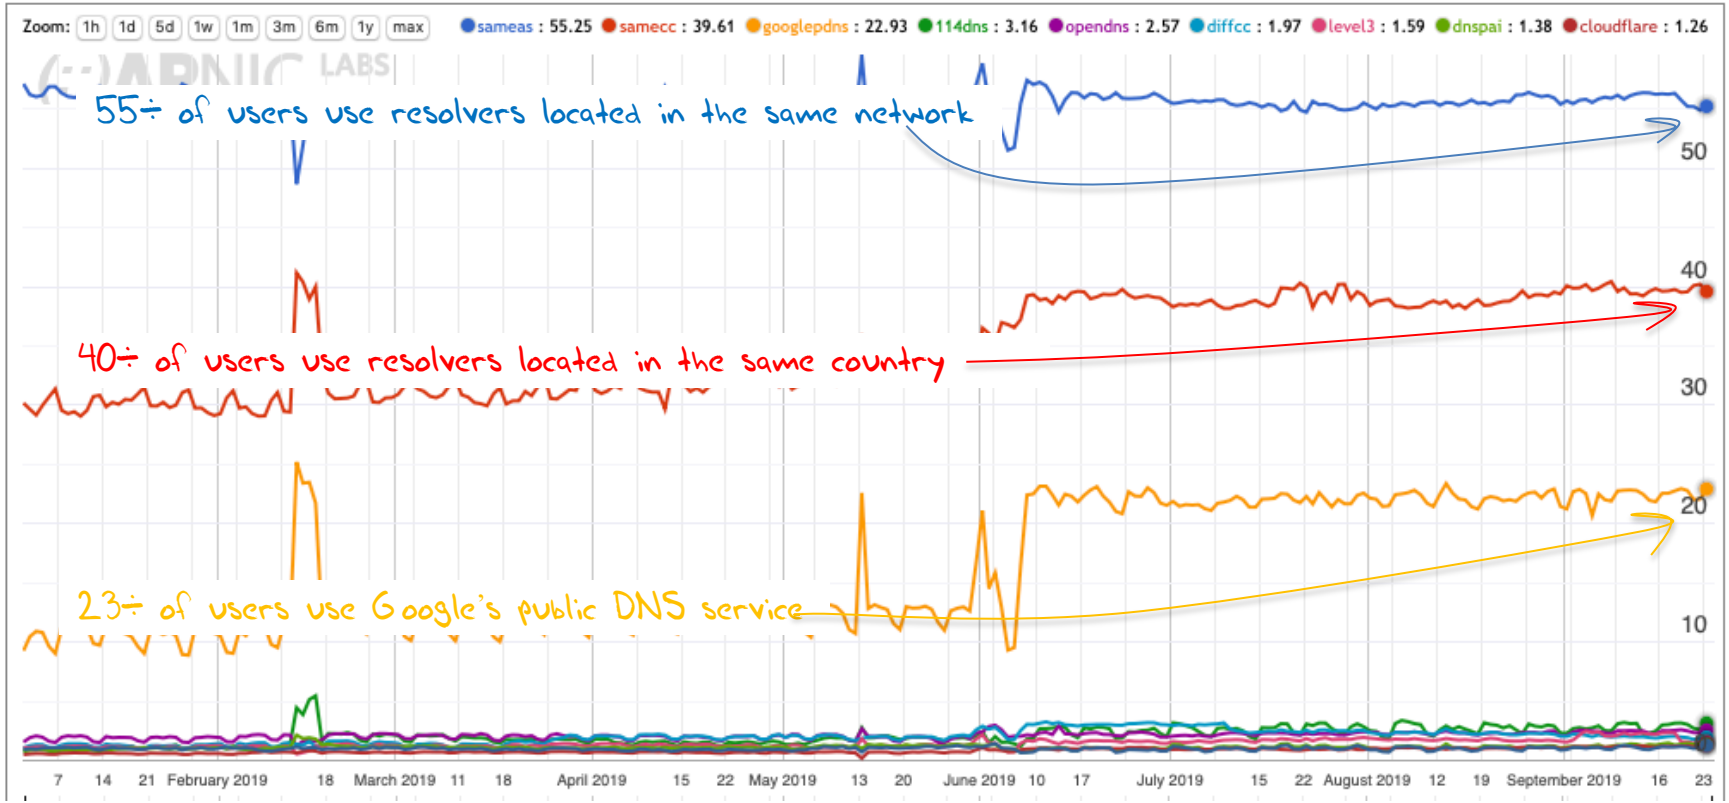

Mapping Open Resolvers

For each country can we show the distribution of the resolvers used by users located within that country?

Where is Google's Public DNS used?

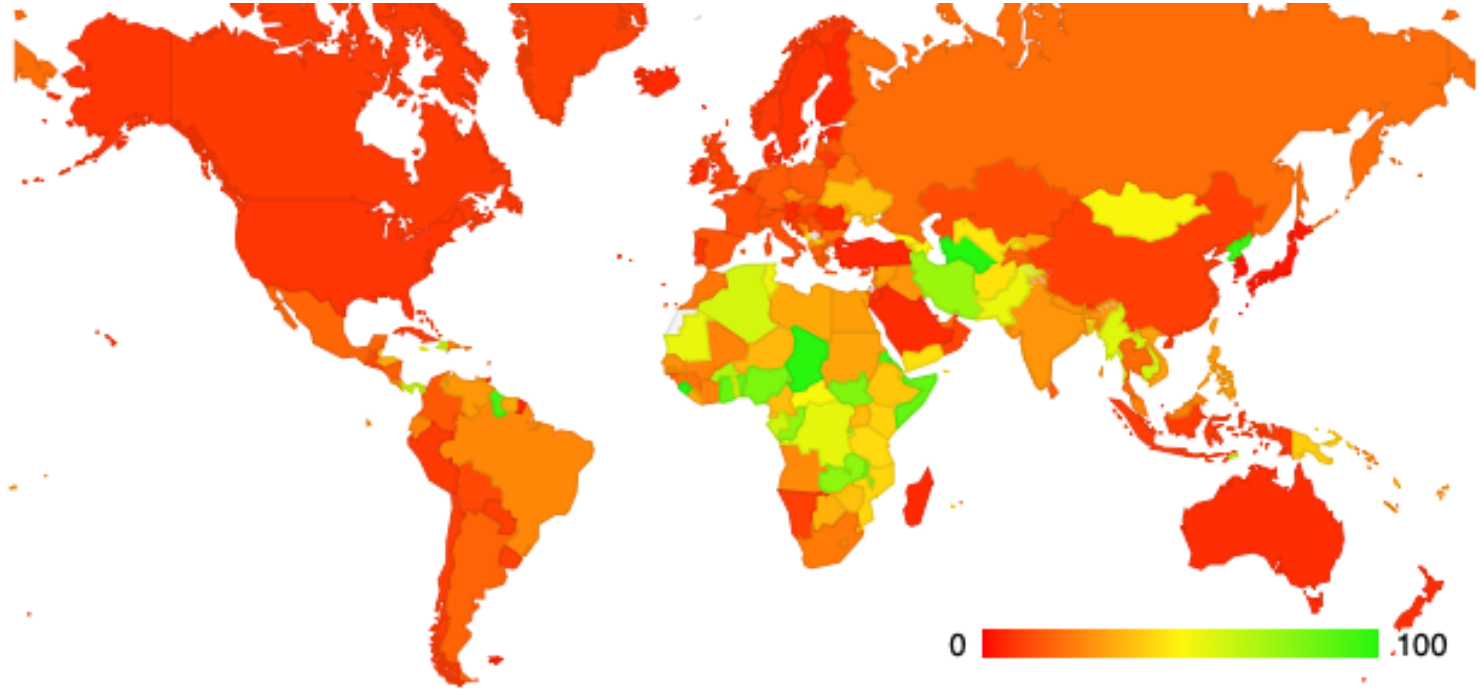

Where are Google's DNS Users?

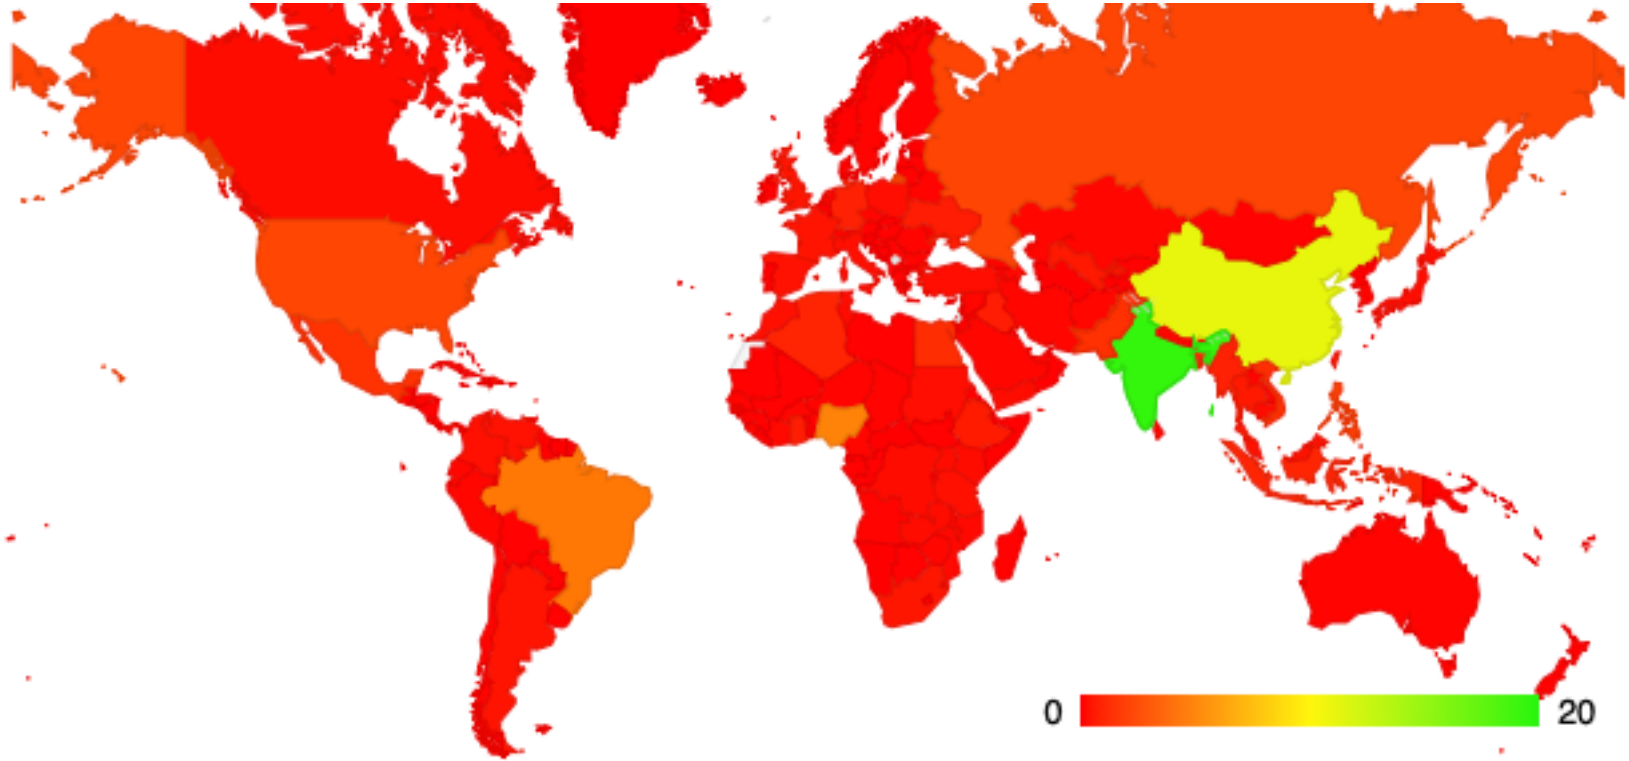

Why is this happening?

- At lot of this story is Google's Public DNS, which now has a "market share" of more than 9% of the Internet's user population for first query
- User's efforts to circumvent content control via national DNS filtering measures
- Network service providers redirecting queries towards Google (Its cheaper than running a local recursive resolver service!)

- <insert your favourite theory here>

Spain?

First Resolver

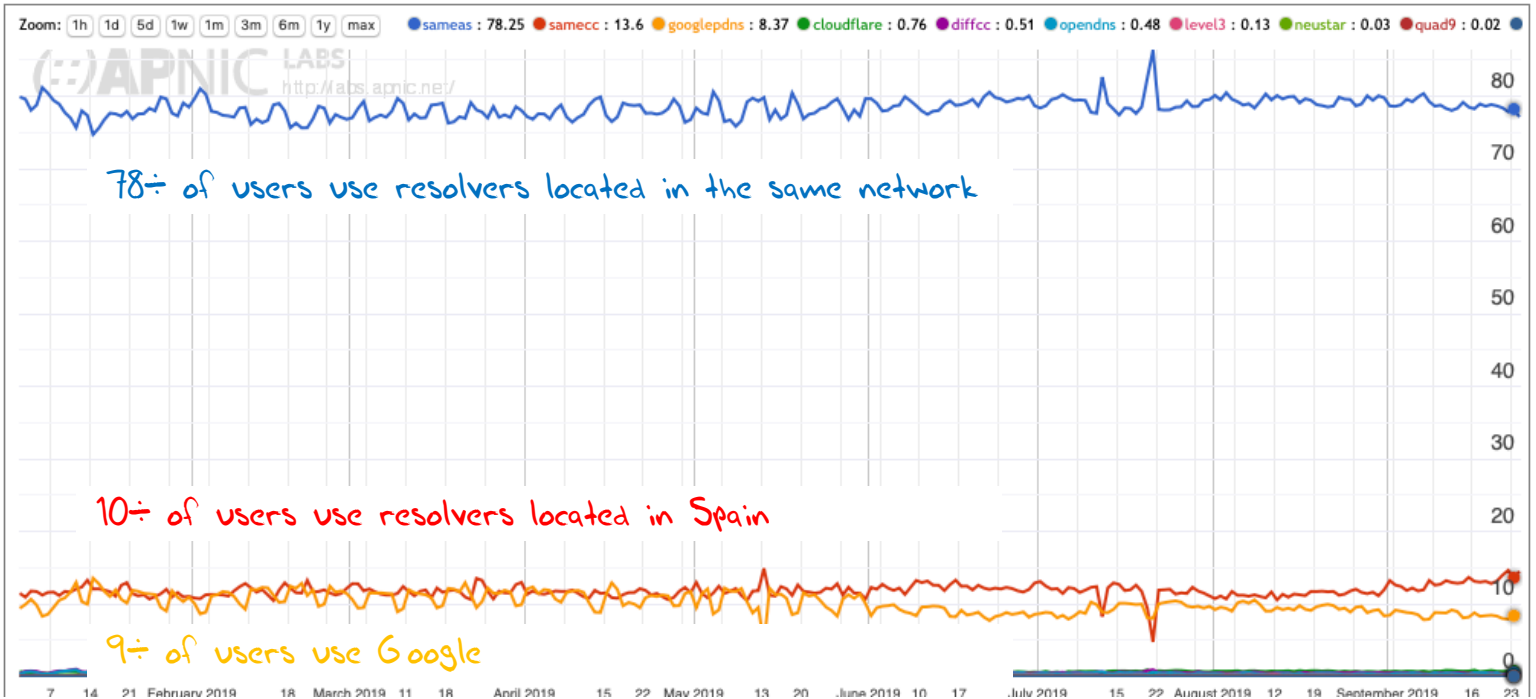

Spain?

All Resolvers

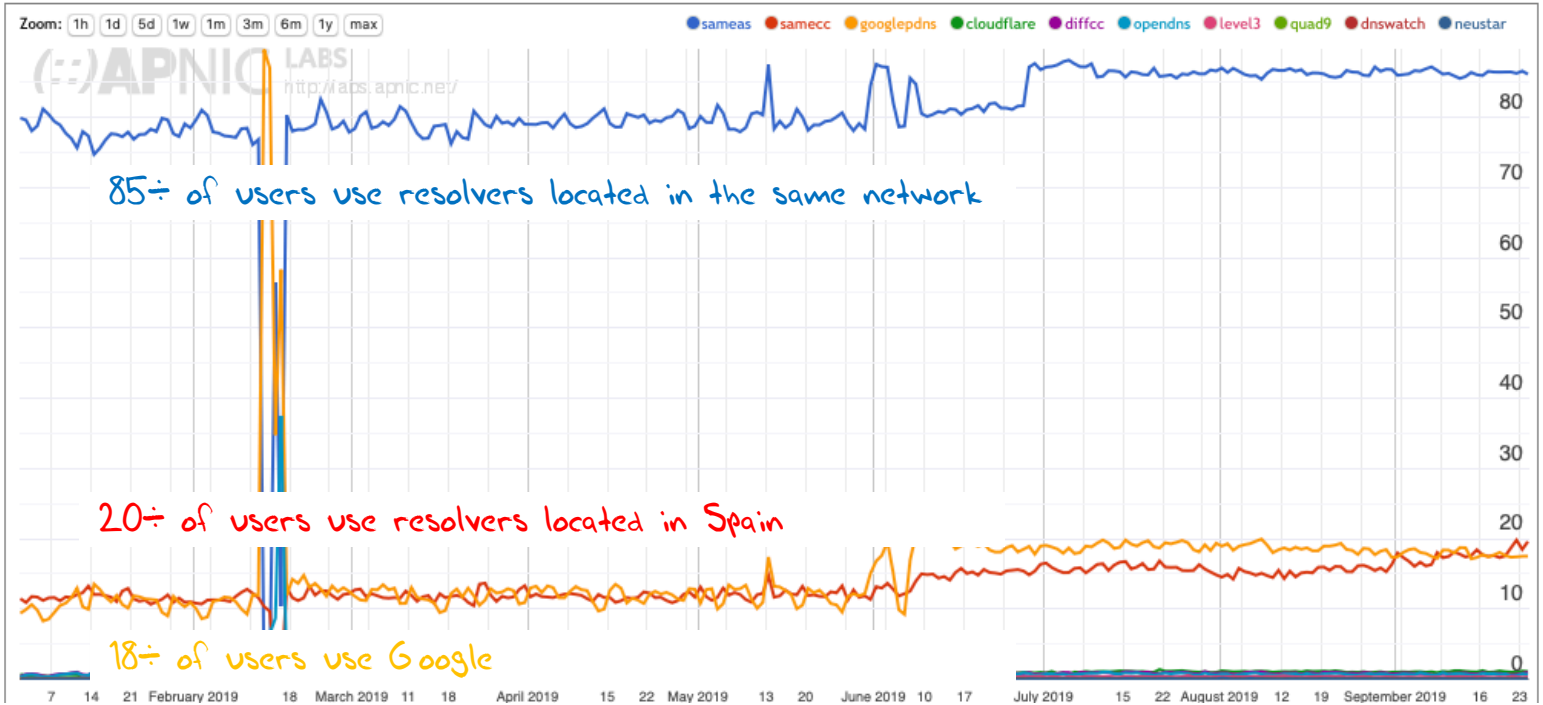

DNS over HTTPS

- Today most DNS is the result of ISP configuration
- Most users are willing to accept the service provider defaults
- But what happens with DNS over HTTPS?
  - What happens when Firefox directs their DNS queries to a different resolver?
  - Will applications exercise greater control of the choice of DNS resolution?

Where is the DNS heading?

- Is the DNS under pressure to aggregate to ever larger resolvers and server farms?
- What is the economic model of name resolution in a highly aggregated environment? Will resolver operators turn to data mining of queries to generate revenue streams?
- Is it possible to reduce the information exposure while still using common resolver caches?
- What is the nature of the trade-off between resolution performance and information leakage in DNS resolution?

Thanks!