The resolvers we use

João Damas

Geoff Huston

As far as I am aware nobody actually said this:

> "In the EU the DNS resolvers we use are dominated by a service operated by an American behemoth that is completely unaccountable to EU users and operates entirely outside of our regulatory framework.
>
> And they are operating an essential piece of Internet infrastructure that many EU Internet users rely on.
>
> We feel that this is unacceptable situation for the EU."

But when I look at the possible motivations behind the DNS4EU initiative it sure feels like this is a credible explanation!

“The resolvers we use”

- What does this ***mean?***
- If you wanted to conduct a measurement experiment to calculate the “market share” of DNS resolvers, then what exactly should we be measuring?

Our Initial thoughts

- Use an Ad to send each user a unique DNS name and look at the authoritative server and collect the IP address of the resolvers asking the recursive resolver, and use Ad data to match this query to an end user IP address

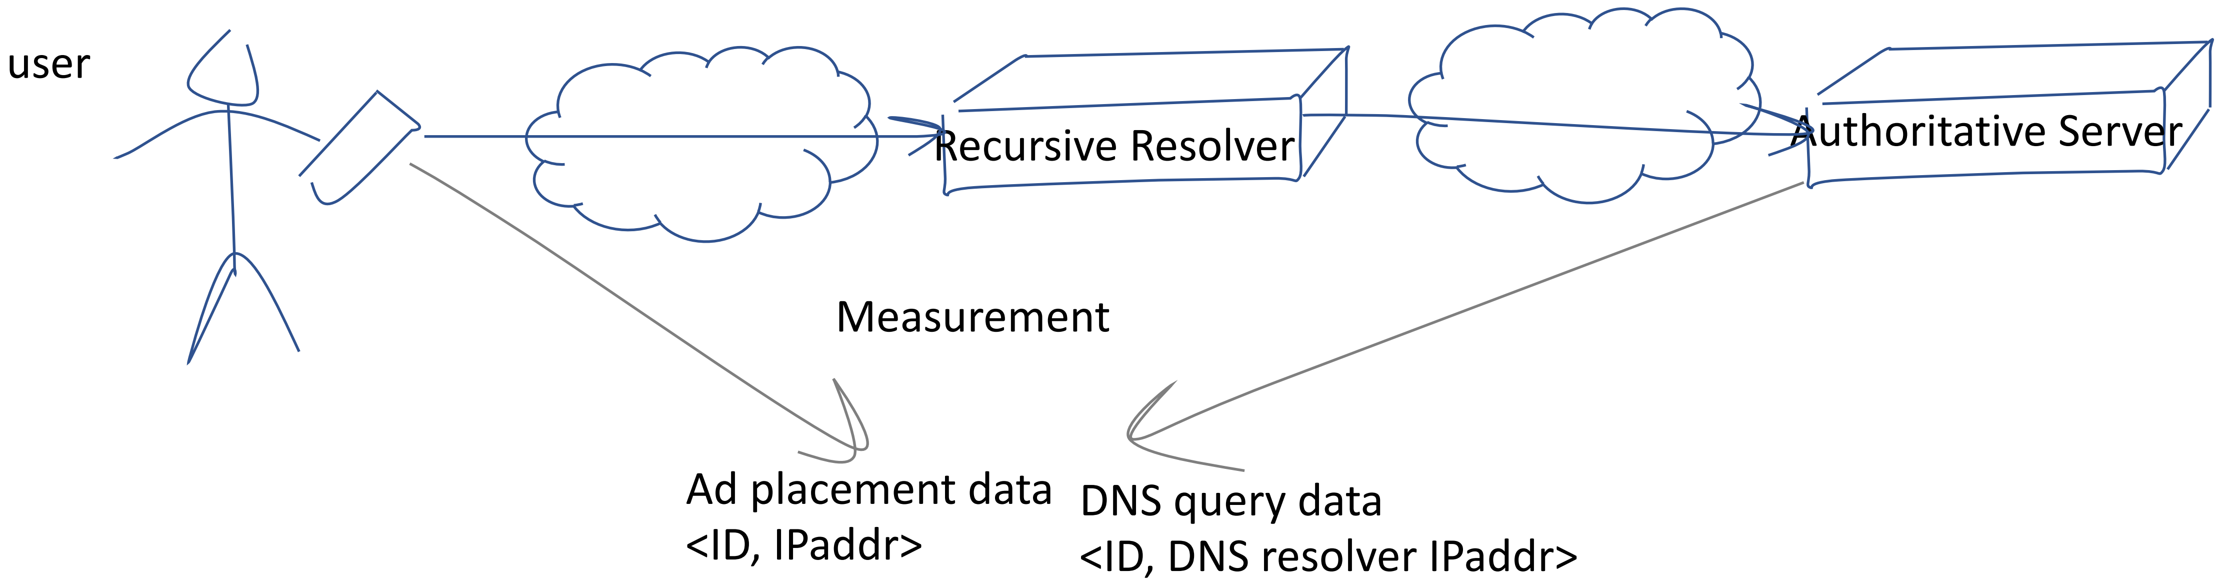

Our Initial thoughts

- Needed to map resolver “helper” addresses to resolver service
  - Which back-end DNS addresses are used by each open resolver?
  - RIPE Atlas helped here for those cases where the open resolver operator does not publish this information
- We map resolvers into a number of categories based on the resolver’s IP address:
  - Resolver is in the same AS as the end user
  - It’s a known Open DNS resolver
  - Resolver is geo-located to the same CC as the end user
  - Resolver is geo-located to a different CC from the end user

DNS in EU

Resolvers seen from a single initial query

Same AS (ISP)	– 70%
Google	– 16%
Same CC	– 8%
Cloudflare	– 5%

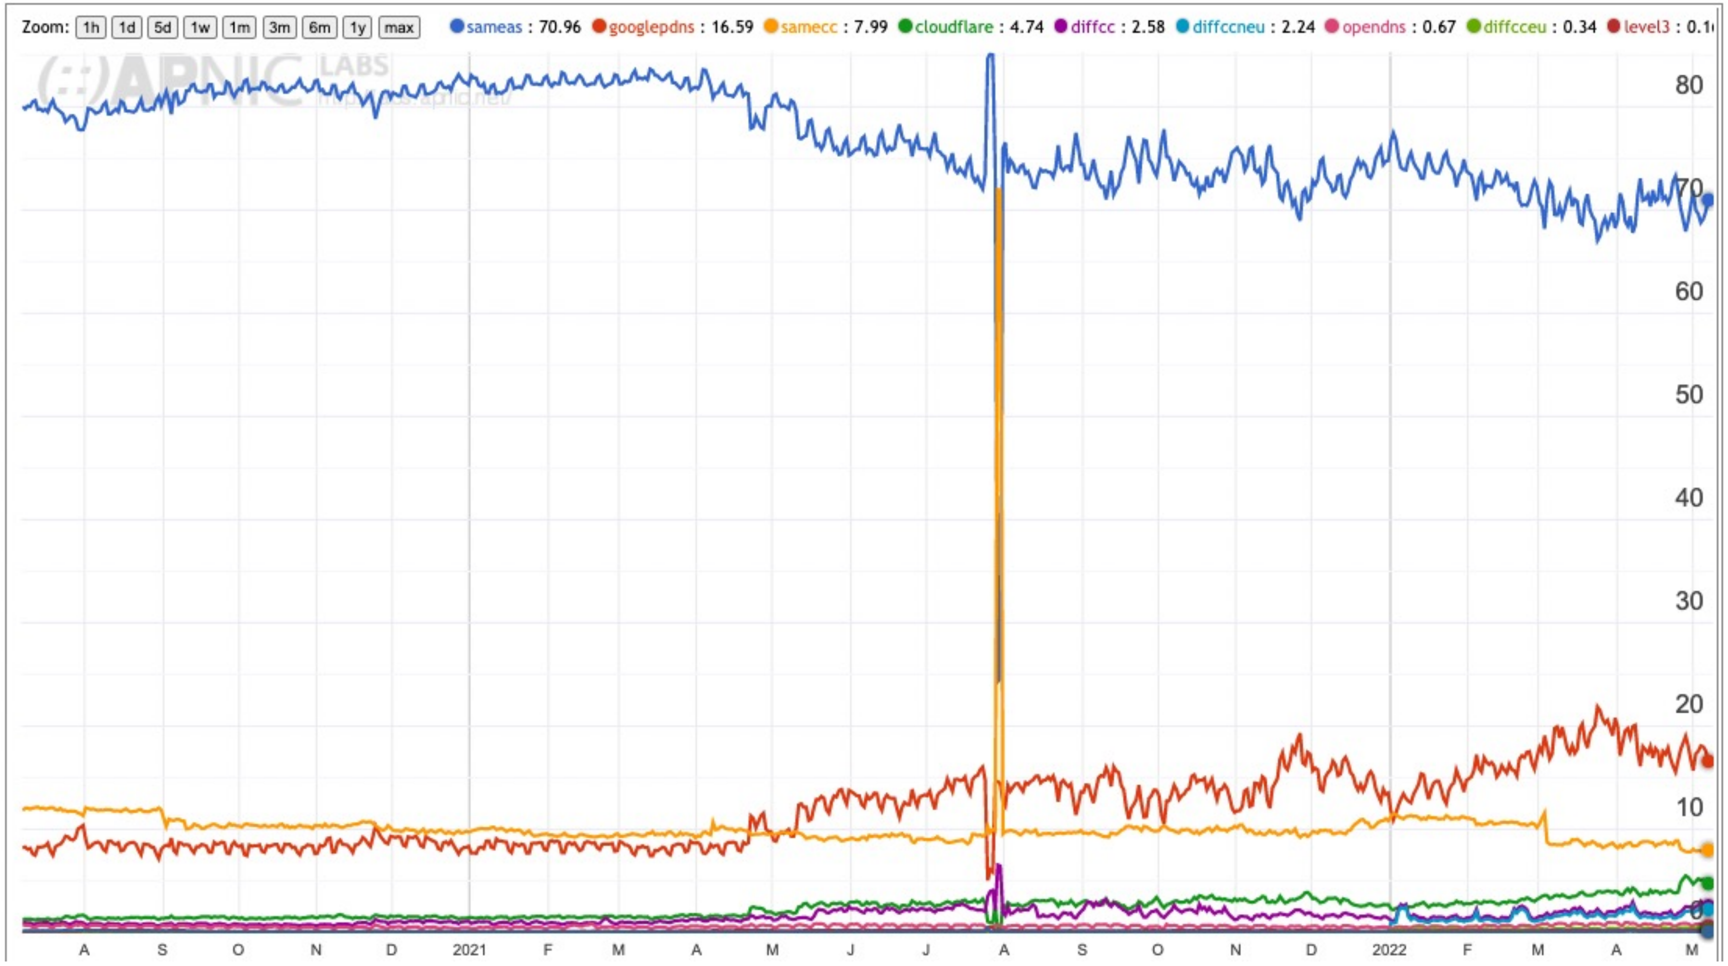

2021

2022

However

- We observe that this single initial query generates a single query to the authoritative server only 73% of the time
- We see an average of 1.65 queries from distinct IP addresses at the authoritative server for each domain name
- What should we do with these “extra” DNS queries?
- In this case we just add them to the count
- Could we do better?

What are we measuring here?

- Seems that this experiment is not clear about what is being measured
- So we thought that maybe we really wanted to know ***all*** the resolvers who ***might*** see your query
- But to flush out all of these resolvers we need to adjust this experiment

Second Thoughts

- Get the authoritative server to return SERVFAIL all the time
- This way the stub resolver is likely to cycle through all the locally configured recursive resolvers to find a non-SERVFAIL DNS response

DNS in EU

Pass 2 Data

All Resolvers seen

from SERVFAIL

Same AS (ISP)	– 72%
Google	– 26%
Same CC	– 9%
Cloudflare	– 6%

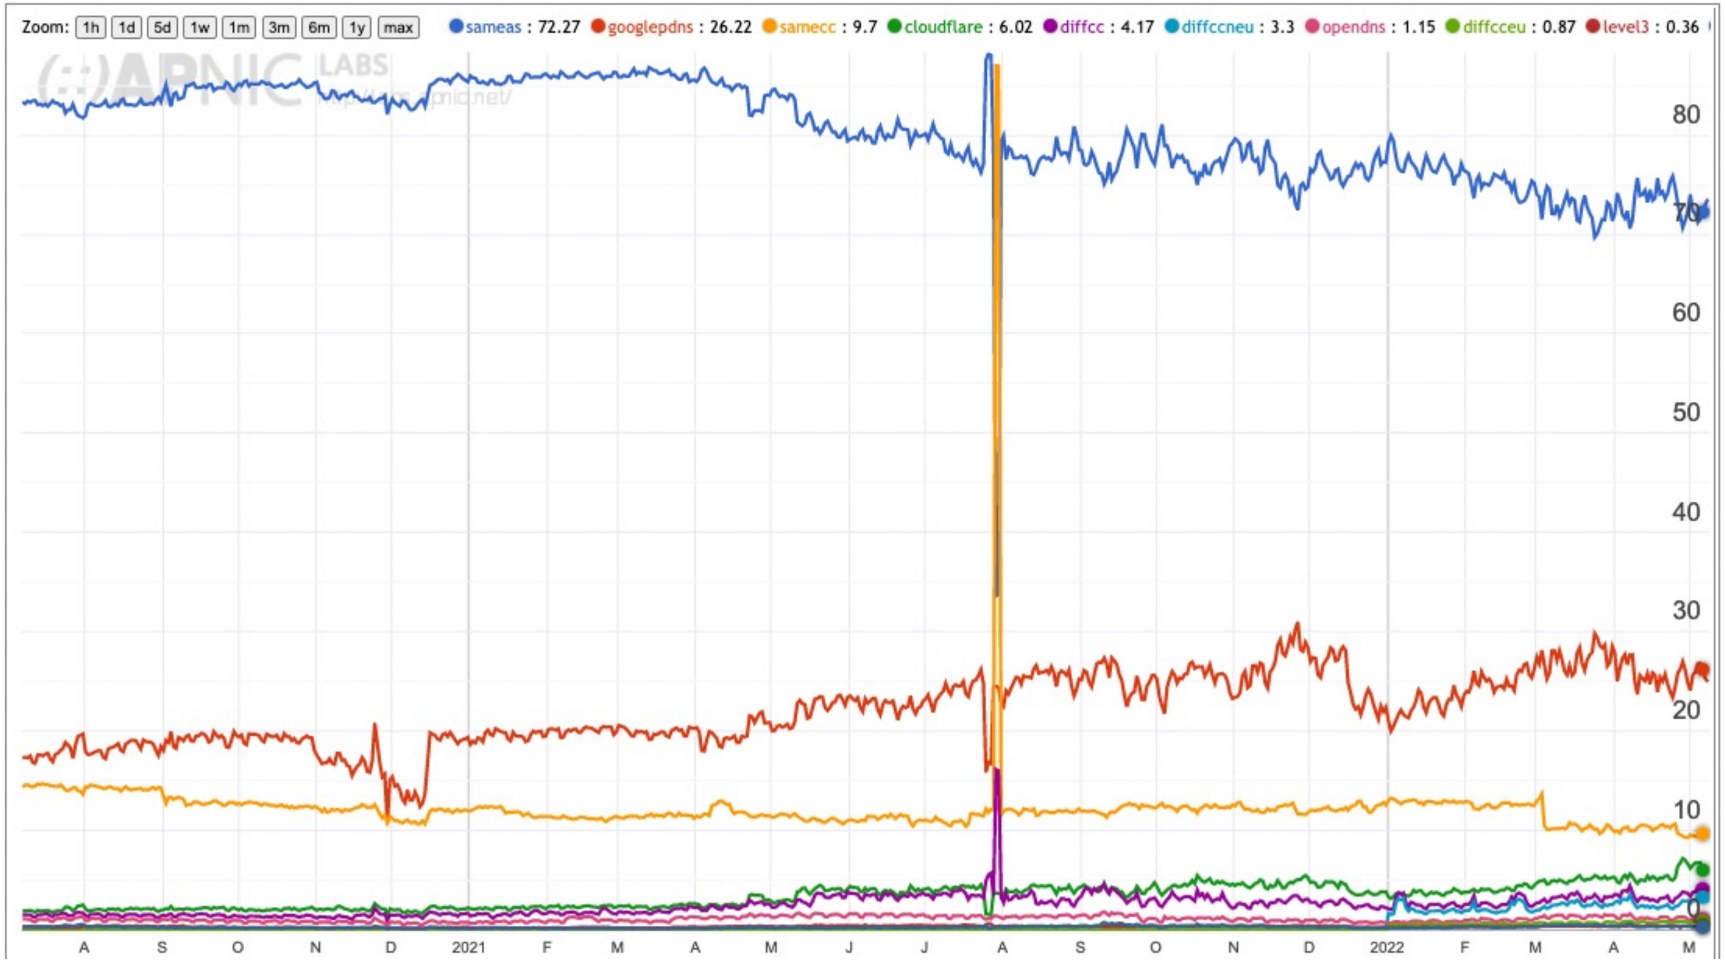

2021 2022

Are we there yet?

- No, not really
  - “might see” your query is not the same as “will see” your query
- Perhaps it is also useful to understand ***which resolver provides the response that the user will use***

Third Pass

- Single query – same as Pass 1
- But only record the first query at the auth server for each unique ID
    - We assume that the first recursive resolver to ask the auth server is the first to provide a response to the stub resolver
- How does this change the measurements?

DNS in EU

Pass 3 Data

First Responder

Same AS (ISP)	– 71%
Google	– 15%
Same CC	– 7%
Cloudflare	– 4%

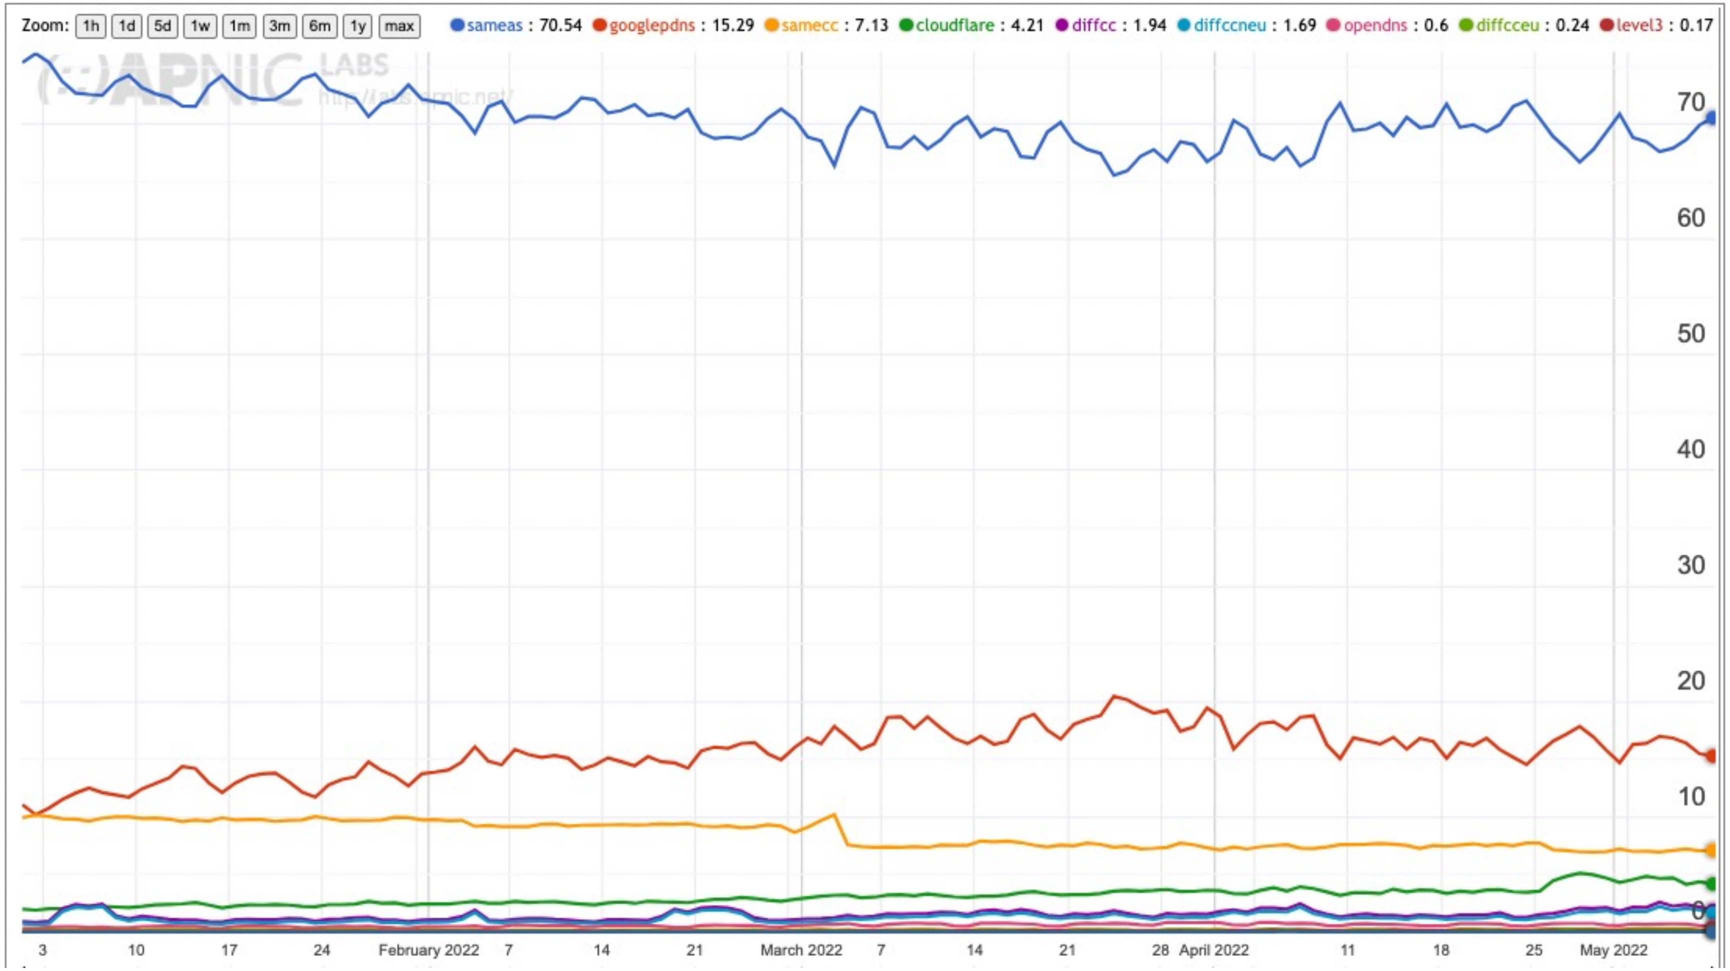

2022

Where are we now?

	All configured Resolvers	Resolvers in the initial query set	First Responder
Same AS	72%	70%	71%
Different AS	9%	8%	7%
Google	26%	16%	15%
Cloudflare	6%	5%	4%

What’s Google’s market share of DNS resolution in EU? 26%? Or 15%

Who is “*we*”?

- Ads are directed to all kinds of end points
  - There are various forms of Enterprise and B2B networks that are measured as well a consumer networks. These enterprise endpoints have a different DNS profile as compared to consumer retail services, as they tend to be more intensive users of open DNS resolver services than we see in consumer networks.
- What if the “we” we are referring to in this measurement are individual consumers in the EU region?
- What happens if we filter the ad data for Europe to look only at consumer retail ISPs?

DNS in EU

Pass 4 Data

First Responder, consumer ISPs

Same AS (ISP)	– 87%
Same CC	– 6%
Google	– 4%
Cloudflare	– 1%

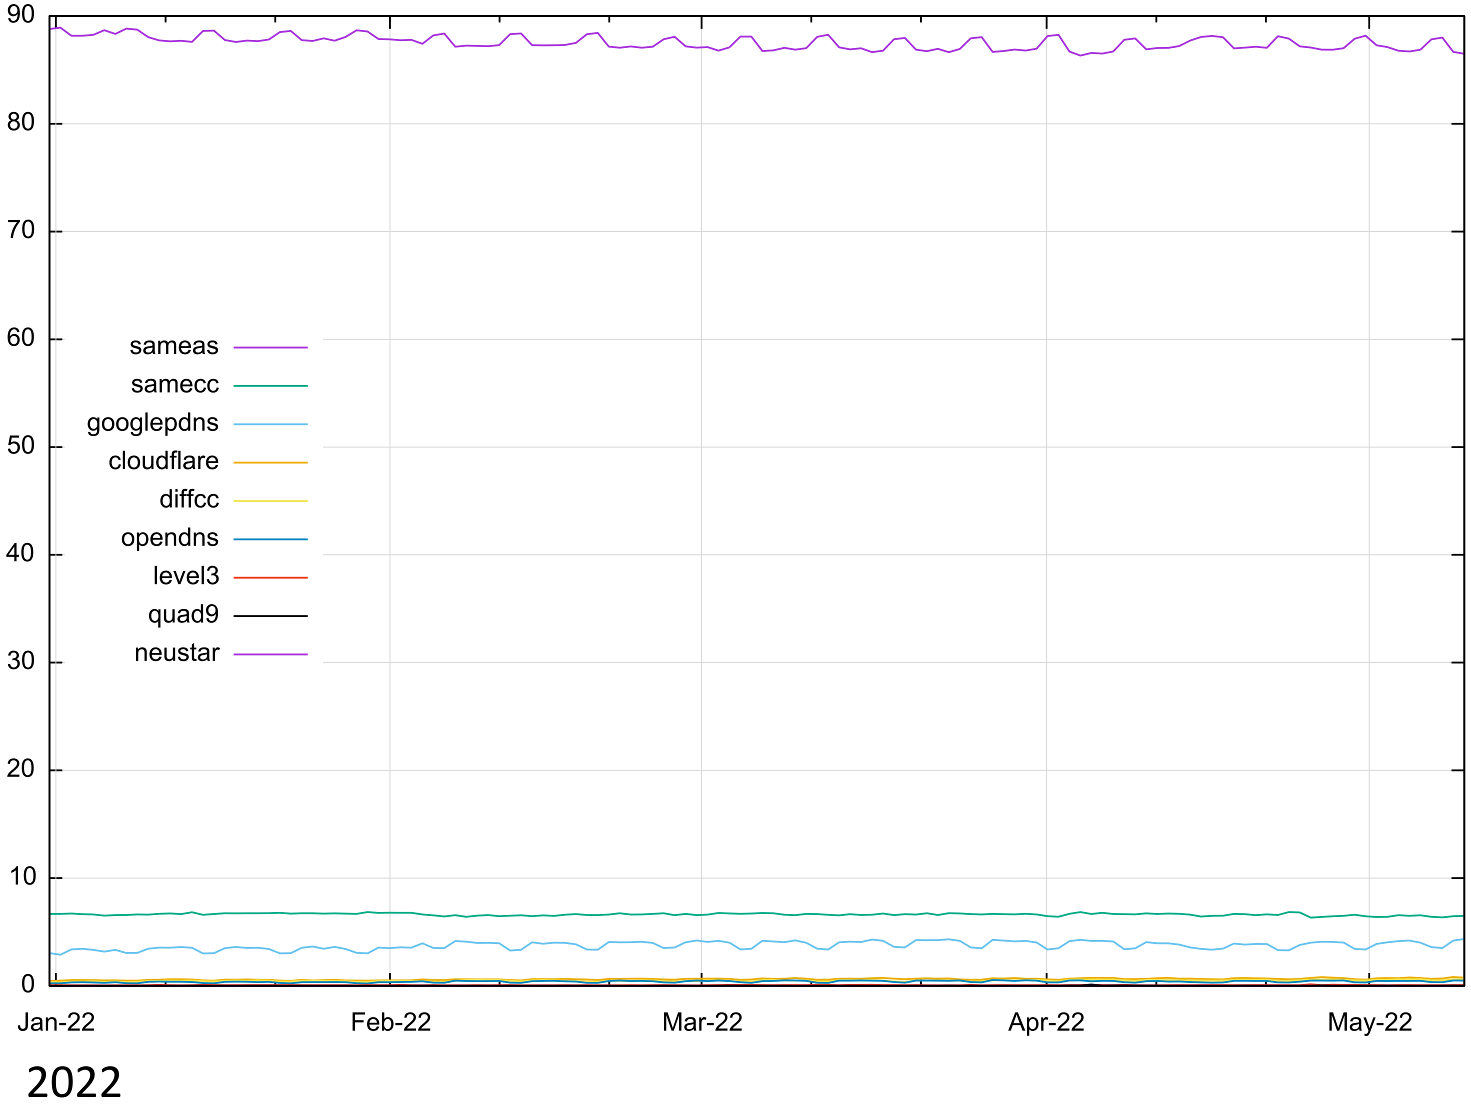

Where are we now?

	All Resolvers	Initial query set	First Responder	Consumer ISPs
Same AS	72%	70%	71%	87%
Different AS	9%	8%	7%	6%
Google	26%	16%	15%	4%
Cloudflare	6%	5%	4%	1%

Most consumers simply follow the ISP provider's default settings

What does all this mean?

- Most users of retail ISPs use their ISP-provided resolver (same AS)
- Some additional users in those ISPs use resolvers in different networks (ASN) but same country
    - This can be due to the ISP having services across different ASNs
- The picture for enterprise networks is slightly different.

So, is there really a problem?

“Yes and No” or perhaps “No and Yes”

- The majority of end users use the ISP-provided default in the EU region
- However, there is an undeniable issue about the emergence of aspects of centrality in the DNS that should concern us all
- So before it becomes a big problem perhaps something can be done?

Can we do something useful in this space?

- Possibly…
- But perhaps not the way the European Comission was proposing to do
- The Internet is at its finest when it self regulates with the network and its users’ interests at the heart of the matter.

OK then…

- Perhaps there is room for the establishment of a common set of operational practices for operators of DNS resolver in all their forms
- RIPE is a good venue to be the seed crystal for this common ruleset
- RIPE has a tried and trusted way to accelerate this initial step
  - RIPE task force