

Measuring the Centralization of DNS Resolution

Geoff Huston, Joao Damas,

APNIC Labs

What are we talking about?

- The DNS is a *highly decentralised* database that distributes its contents over much of the Internet
- The DNS data model also includes information replication (secondary authoritative servers) that attempts to provide resiliency and scalability by removing critical single choke points within the database
- The DNS name resolution protocol includes query fallback to increase the robustness of name resolution
- It all sounds as if the DNS highly diverse and extensively decentralised.

What are we talking about?

- The DNS is a *highly decentralised* database that distributes its contents over much of the Internet
- The DNS data model also includes zone file replication (secondary authoritative servers) to provide resiliency and scalability by removing single choke points within the database
- The DNS name resolution protocol includes query fallback to increase the robustness of the centralisation
- It all sounds like the DNS highly diverse and extensively decentralised.

But is it really decentralised?

Measuring Centrality

- Various measures have been used in the related space of market dominance which appear to have some relevance to the study of market dominance in the DNS
 - Australia's Consumer and Competition agency uses a metric of 70% **market share** by a single entity
 - Or there is the four-firm **concentration ratio** which uses the market share of the four largest firms
 - Or there is the **Hirfindahl-Hirschman index**

Herfindahl-Hirschman Index

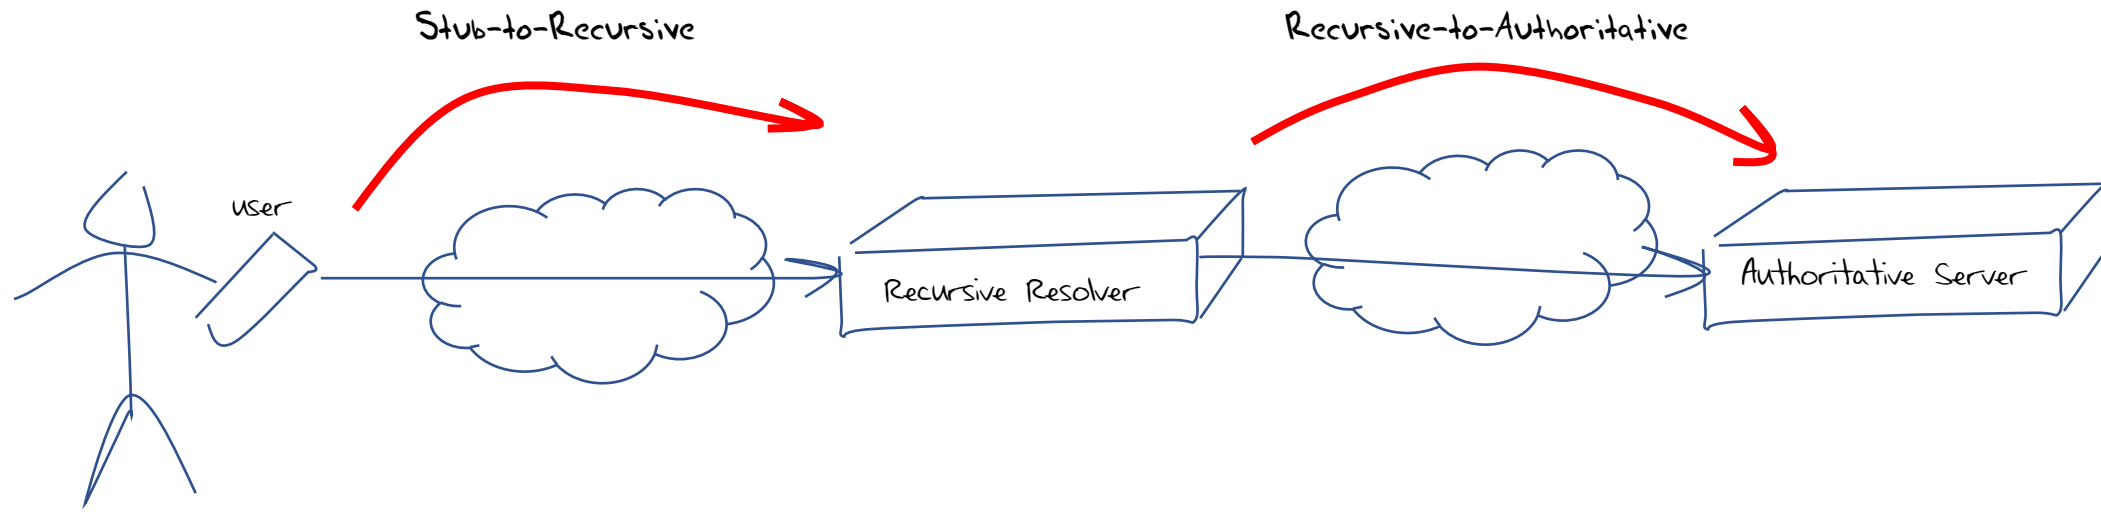
- The HHI is used in market analysis to indicate the level of competition between market entities. It is the average market share of the market, weighted by market share
- Sum of the square of the market share (%) of the top 50 entities
- Above 25% is often taken as an indicator of market skew
- Above 10% is would be considered as a market showing “moderate concentration”

The DNS resolution environment

- The question of **centrality** in the DNS resolution environment is equivalent to the questions of **market dominance** and **market concentration**
- So lets look at DNS resolution as a market and use these measurements to assess the degree of concentration in the supply of DNS resolution services

Looking at concentration in DNS Name Resolution

- I. How centralized is the **recursive resolver function**?
- II. How centralized is the **authoritative server function**?




I Recursive Resolution

Measuring Recursive Resolution

We use Ads to send each user a unique DNS name to resolve.

We use an authoritative server as the data collection point and collect the IP address of the resolvers asking the authoritative server

We then use the Ad presentation data to match this query to an end user IP address




Tuning the measurement

- The Authoritative Server always answers queries immediately with the A / AAAA records as requested
- The data is unsigned and the responses fit comfortably within 512 octets of DNS payload
- We try to minimise timeouts and requeries by steering users to a DNS Authoritative server that is (roughly) on the same continent as the user

Mapping

- We need to map the resolver “helper” addresses to a 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. The categories we use:
 - 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

Results



Same AS


Google
Same CC

Cloudflare
Others

June 2020

October 2022

Results




- Two-thirds of users direct their queries to the recursive resolver that is operated by their ISP
- One seventh (15%) of users have their queries resolved by Google's Public DNS resolver
- One seventh (15%) of users direct their queries to a recursive resolver that is geolocated to the same country as they are – probably their ISP using a resolver in a different AS
- Everything else – nothing more than 3%

- Is the recursive resolver market centralized? Probably not
- Is the **open** recursive resolver market centralized? That's a different question!
- But first let's understand what is being measured here

However, the measurement is not as simple as it may suggest

- We observe that this single initial query generates 1 or more queries from a single recursive resolver IP address just 30% of the time
- 2 or more different resolvers are queried in 60% of cases
- Most of the time (90% of cases) these multiple resolver IP addresses are all in the same AS

Cumulative Distribution of number of resolver IP addresses seen to query for a unique DNS name



Multiple resolvers "see" individual stub queries

- We see an average of 3.23 distinct resolver IP addresses at the authoritative server for each queried domain name within the first 15 seconds
- What should we do with these "extra" DNS queries?
- In this case we just add them to the count
- So we are measuring who "sees" my DNS queries


What are we measuring here?

- 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

Seeing Everything!

- 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

All Recursive Resolvers




- Two-thirds of users direct their queries to the recursive resolver that is operated by their ISP
- ~~One seventh (15%)~~ fifth (20%) of users have their queries resolved by Google's Public DNS resolver
- ~~One seventh (15%)~~ sixth (16%) of users direct their queries to a recursive resolver that is geolocated to the same country as they are – probably their ISP using a resolver in a different AS
- Everything else – nothing more than ~~4%~~ 3%

How many resolvers see the query now?

- We observe that this single initial query generates 1 or more queries from a single recursive resolver IP address just 12% of the time
- 2 or more different resolvers are queried in 30% of cases
- Most of the time (75% of cases) these multiple resolver IP addresses are all in the same AS

Cumulative Distribution of number of resolver IP addresses seen to query for a unique DNS name when the response is SERVFAIL




Are we there yet?

- No, not really
- 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?

First Query Results



- Two-thirds of users direct their queries to the recursive resolver that is operated by their ISP
- One ~~fifth (20%)~~ ^{seventh (15%)} of users have their queries resolved by Google's Public DNS resolver
- One ~~sixth (18%)~~ ^{eighth (1%)} of users direct their queries to a recursive resolver that is geolocated to the same country as they are – probably their ISP using a resolver in a different AS
- Everything else – nothing more than ~~4%~~ ^{2%}

What are we looking at?

- Who **gets** to see my queries?
- Who **might** get to see my queries?
- Who do I **believe** for answers?

Concentration Measurements

- Lets look at the “market” of DNS open resolvers using the “all resolvers” measurements
 - Single Entity Dominance:
 - Google has 68.7% of the open DNS resolver market
 - Four-Firm Concentration:
 - Google, Cloudflare, 114DNS and OpenDNS have 91.6% market share
 - HHI Index:
 - 49%
- So the open resolver market sector is **highly centralized**.
- But this open resolver activity represents only one third of the total resolution market, and the HHI Index of the open resolvers as a subset of the total resolution market is far lower, at 5%

II Authoritative Servers

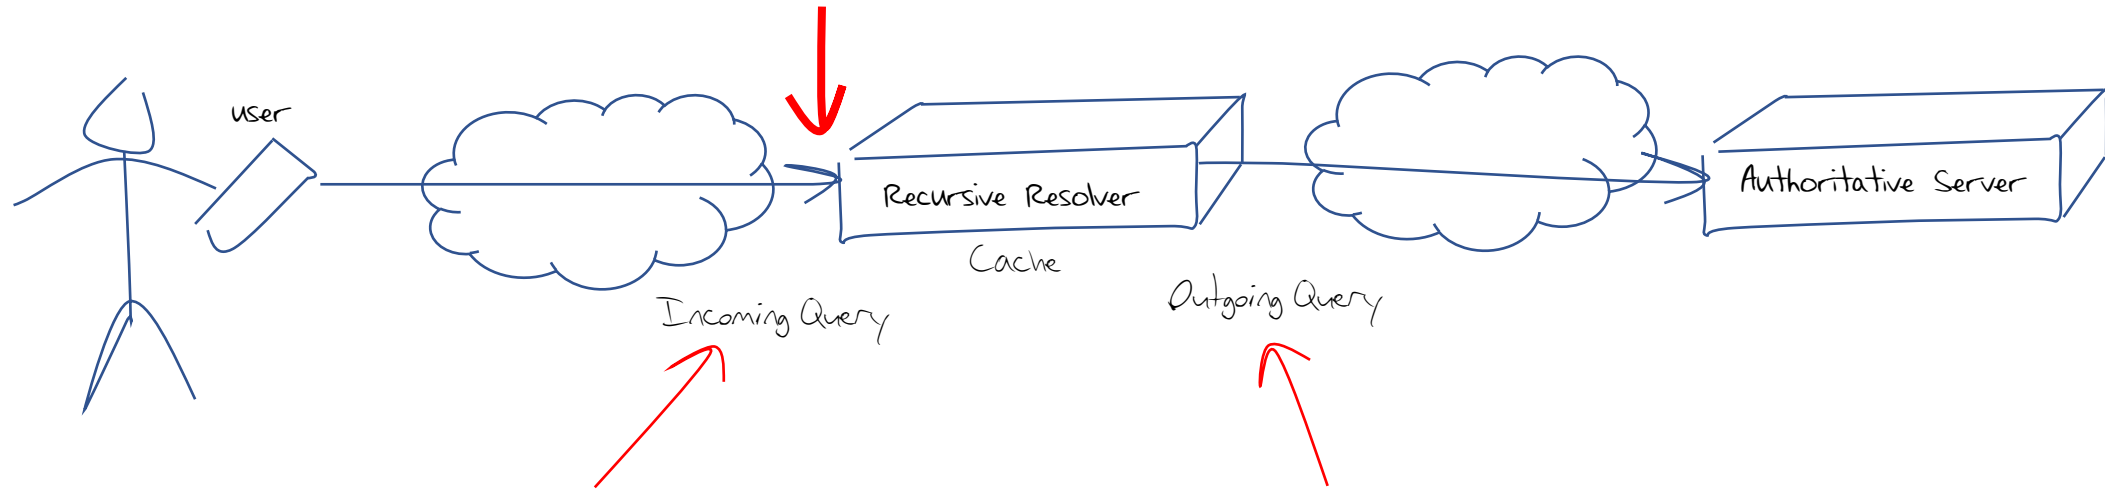
- Data about the recursive-to-authoritative query set is hard to find
 - Recursive resolvers sit in a privileged position in the DNS as they are exposed to both the identity of the stub resolver (the ‘user’) and the DNS names that they are querying
 - So there are many caveats that apply to access to such data – and rightly so
- At APNIC we have limited access to the data relating to the use of the 1.1.1.1 recursive resolver
 - We don’t know who is querying, but we can see query names and query protocol
 - The market share of Cloudflare’s open resolver service is around 3% of users which is a non-trivial resolver in the open resolver set (ranked #2 in terms of market share)

Centrality in Authoritative Servers?

- One way to measure this is to look at the **query-count weighted ranking** of the DNS authoritative server providers
- If an authoritative name server hosts a very popular domain name then it's likely that the query count will be high
- If a service operator hosts a large number of domains on its authoritative server infrastructure, then it's likely that the query count will be high
- So lets characterise the authoritative service hosting market by their query-based 'market share'

What's a "query"?

We use incoming queries to determine relative weight



The query count at this point depends on stub activity rather than cache settings

The query count at this point is dependant on the cache settings

Measurement Technique

- Obtain a data set of 24 hours of query name data from the 1.1.1.1 resolver system
- Group the query names
- Resolve the names to find the “lowest” authoritative name server for the query name using a local resolution environment
- Take the first name server name
- Discard the query names
- Resolve the name server names to IP address, and discard the name server names
- Map the IP addresses to AS numbers, and discard the IP addresses
- Group the query counts into AS numbers and rank by query share

Data Set


Rank	AS Number	Query Share	Cumulative Share	Name
56	31034	0.06%	89.64%	Aruba, IT
57	20446	0.06%	89.70%	Stackpath CDN, US
58	199524	0.06%	89.76%	Gcore, LU
59	60068	0.05%	89.82%	CDN77, GB

Here's an extract of the resultant data set


A 24 hour data capture identified 26,971 unique AS numbers (out of a total of 75,000 unique AS numbers in the routing table)

While approximately one third of networks host at least one queried authoritative name server the top 50 ASNs have 89.2% of the query share.

Cumulative Distribution



Cumulative Distribution



Top 10 Auth Server Networks

Rank	AS Number	Query Share	Cumulative Share	Name
1	AS16509	35.7%	35.7%	Amazon-02, US
2	AS13335	9.3%	45.0%	Cloudflare, US
3	AS15169	8.3%	53.3%	Google, US
4	AS21342	4.0%	57.3%	Akamai – ASN2, US
5	AS8068	3.9%	61.2%	Microsoft, US
6	AS397239	3.7%	64.9%	UltraDNS, US
7	AS714	3.4%	68.3%	Apple, US
8	AS31898	3.1%	71.4%	Oracle, US
9	AS*	2.5%	73.9%	Root Server System
10	AS62597	2.5%	76.4%	NSONE, US

Concentration Measurements

- Lets look at the “market” of DNS authoritative server providers using query-weighted ranking
 - Single Entity Dominance:
 - Amazon has 35.7% of the Authoritative Server market
 - Four-Firm Concentration:
 - Amazon, Cloudflare, Google, and Akamai have 57.3% market share
 - HHI Index:
 - 15%
- This appears to be a “**moderately concentrated**” market

Side note

- What's that root query volume?
 - The root servers report that around 2/3 of the queries seen at the root result in NXDOMAIN responses
 - (Duane Wessels, DNS OARC?? 20??)
 - But this data indicates that only 2.5% of queries seen by this recursive resolver are NXDOMAIN queries
- Why?

Geopolitical Centrality?

- There are 10 network entities who host the authoritative name servers that have a query share of three quarters of the recursive-to-authoritative DNS query volume
- All 10 networks are attributed to US entities

Caveats

- The query sample set is not completely uniform and there is a potential bias to enterprise use and some browser use
- Using query volumes as a proxy for some form of market share is not a universally accepted analytic metric

Thanks!