Root DNS Anycast in UK & Ireland

INEX June Meeting, Dublin

What are root DNS servers?

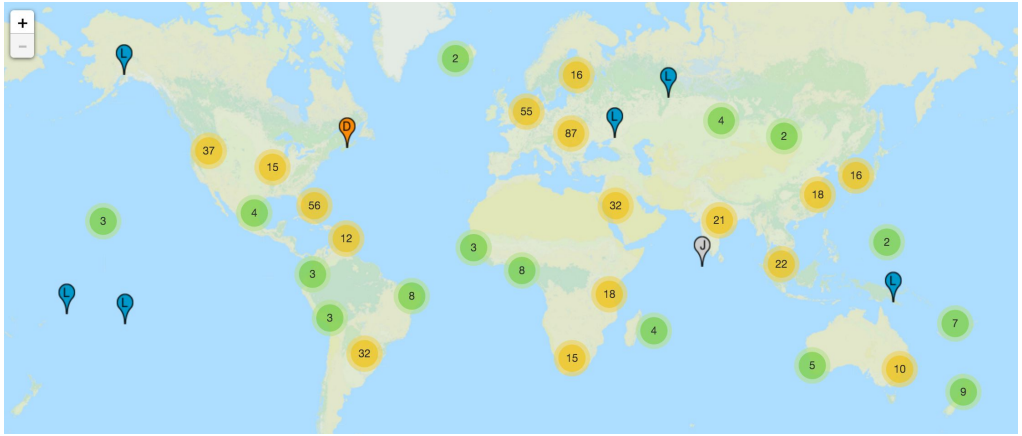

- Authoritative DNS servers for top level “dot”
- Knows authoritative DNS server of each TLD & it’s glue
- Logically 13 servers (from a to m)
- Heavily anycast across hundreds of servers located across the world

What is anycast?

- Announcing same address pool from multiple locations
- In theory routers sending traffic off to nearest anycast node
- Works well with lot of peering but limited set of transits
- Breaks often due to networks having own preference for each other networks

Testing Methodology

Two ways to test anycast performance:

1. Triggering dig CHAOS id.server @$ROOT txt from root servers

2. Triggering ping to the root servers from RIPE Atlas probes in the region

Testing Methodology

Two ways to test anycast performance:

1. Triggering dig CHAOS id.server @$ROOT txt from - *Works for 9 out of 13 root DNS servers*

2. Triggering dig CHAOS hostname.bind @$ROOT txt from - *Works for 4 out of 13 root DNS servers*

3. Triggering ping to the root servers from RIPE Atlas probe in the region - *Works for 12 out of 13 root servers*

Limitations

1. No way to practically test IPv6 due to very low number of IPv6 enabled RIPE atlas probes.

2. Due to blocked ICMP - G root was excluded from comparison of latency.

3. Diversification of RIPE Atlas probes across ASNs is not well tested. It is there in some cases while missing in other cases.

4. A considerable number of RIPE Atlas probes are on DSL and it may have it’s own overhead in the resolution.

Root servers which support CHAOS Class Queries

Root Server	id.server support	hostname.bind support
a.root-servers.net.	Yes	No
b.root-servers.net.	No	Yes
c.root-servers.net.	No	Yes
d.root-servers.net.	Yes	No
e.root-servers.net.	Yes	No
f.root-servers.net.	Yes	No
g.root-servers.net.	No	Yes
h.root-servers.net.	Yes	No
i.root-servers.net.	Yes	No
j.root-servers.net.	Yes	No
k.root-servers.net.	Yes	No
l.root-servers.net.	Yes	No
m.root-servers.net.	No	Yes

Countries with root DNS servers

1. UK (13)
2. Ireland (3)

Countries with active RIPE Atlas Probes

1. UK (611)
2. Ireland (109)

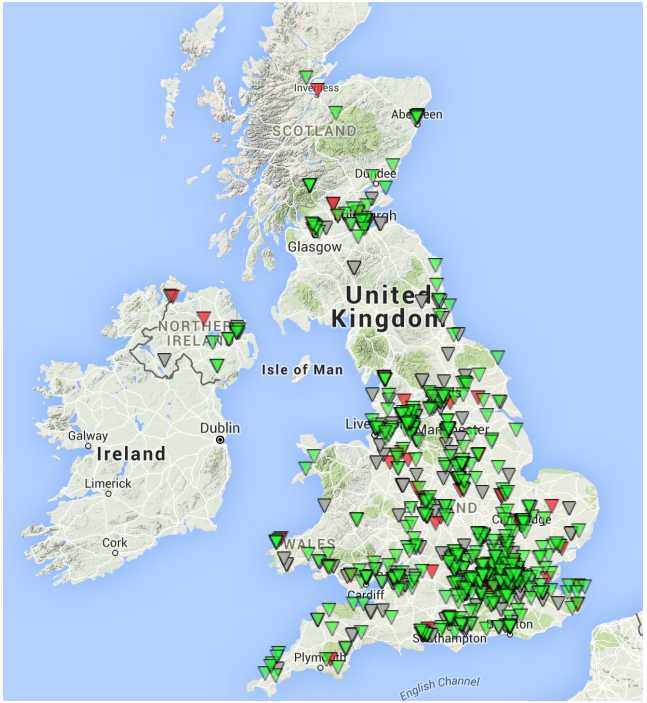

Root DNS Anycast in UK

1. A root - London
2. D root - London
3. D root - Leeds
4. D root - Manchester
5. E root - London
6. F root - London
7. I root - London
8. J root - Leeds
9. K root - London
10. L root - London
11. L root - Leeds
12. L root - Rochester
13. L root - Dundee

Root DNS Anycast in UK

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	Yes	12.9	100.00%	0.00%	0.00%
B root	No	143.2	0.00%	0.00%	100.00%
C root	No	17.93	0.00%	100.00%	0.00%
D root	Yes	22.75	62.22%	20.00%	17.78%
E root	Yes	37	95.12%	4.88%	0%
F root	Yes	28.39	36.17%	34.04%	29.79%
G root	No		0.00%	82.61%	17.39%
H root	No	93.64	0.00%	100.00%	0.00%
I root	Yes	53.05	55.32%	21.28%	23.40%
J root	Yes	69.5	0.00%	51.06%	48.94%
K root	Yes	39.42	63.04%	19.57%	17.39%
L root	Yes	57.88	15.22%	54.35%	30.43%
M root	No	37.04	0.00%	93.62%	6.38%

Root DNS Anycast in UK

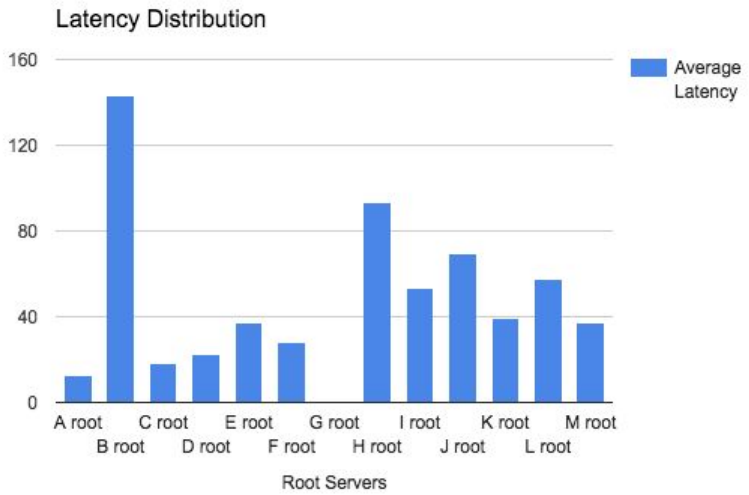

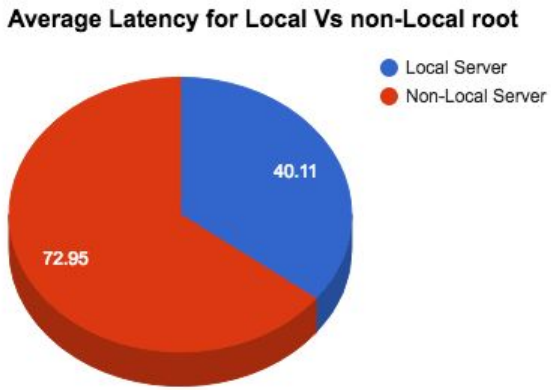

Root DNS Anycast in Ireland

1. D root - Dublin
2. E root - Dublin
3. J root - Dublin

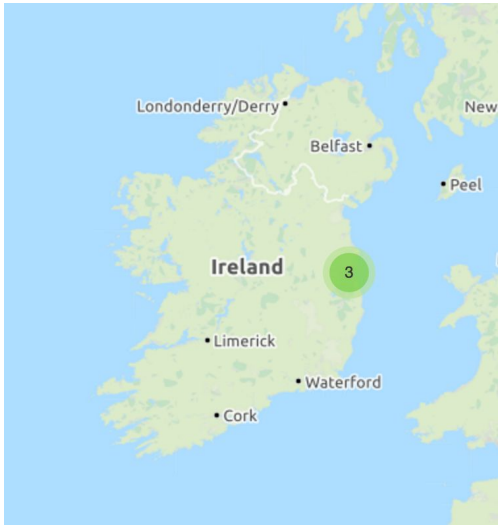

Root DNS Anycast in Ireland

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	No	29.26	0.00%	100.00%	0.00%
B root	No	178.24	0.00%	0.00%	100.00%
C root	No	71.14	0.00%	57.45%	42.55%
D root	Yes	26.47	97.87%	0.00%	2.13%
E root	Yes	22.75	97.87%	2.13%	0.00%
F root	No	35.63	0.00%	89.36%	10.64%
G root	No		0.00%	93.62%	6.38%
H root	No	117.1	0.00%	0.00%	100.00%
I root	No	48.99	0.00%	89.36%	10.64%
J root	Yes	17.9	97.87%	2.13%	0.00%
K root	No	38.44	0.00%	97.87%	2.13%
L root	No	44.21	0.00%	97.87%	2.13%
M root	No	31.76	0.00%	100.00%	0.00%

Root DNS Anycast in Ireland

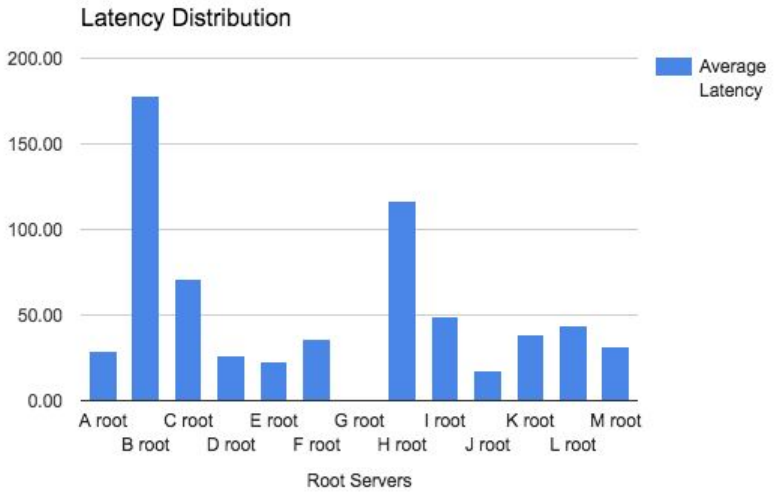

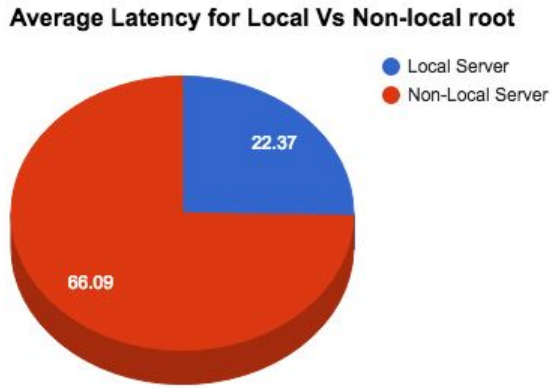

Why does anycast fail?

1. Networks often prefer a customer route over peering, over transit path.
2. BGP best path is not always the geographically best path.
3. Things break due to route leaks where announcement propagates beyond geography.
4. Use of “no-export” in peered routes is debatable.

Misc Points about the study

1. End users “speak to” DNS recursor of ISP and not root DNS servers directly.

2. Presence of even a single root server impacts as DNS recursor software picks it up based on performance & hence low latency with just one or more root DNS server helps in overall resolution time.

3. Users not using local ISPs server & relying on popular open DNS recursors have different resolution path alltogether.

Conclusions

1. More root DNS servers are good for a country, it reduces latency considerably.

2. IXP where root DNS can peer with large number of networks is good.

3. Apart from latency, more DNS servers ensure low impact of submarine or long distance terrestrial cable cuts.

4. It’s better to buy IP transit from a provider with a network & peering in large geography as compared to a localized player.

References

- RIPE Atlas Project - http://atlas.ripe.net
- Root DNS servers Information - http://root-servers.org/
- DNS Chaos Class (Chaosnet) - https://en.wikipedia.org/wiki/Chaosnet
- Root DNS Zone - ftp://ftp.rs.internic.net/domain/root.zone
- Anycast - https://en.wikipedia.org/wiki/Anycast
- Kabindra Shrestha from PCH - https://www.pch.net

Thankyou!

Questions?

Peering?

Email: anurag@he.net
Twitter: @anurag_bhatia

ASN: 6939
Web: http://he.net
Peeringdb: http://as6939.peeringdb.net