Root DNS Anycast in South Asia

Anurag Bhatia, Network Researcher
JANOG 38

What are root DNS servers?

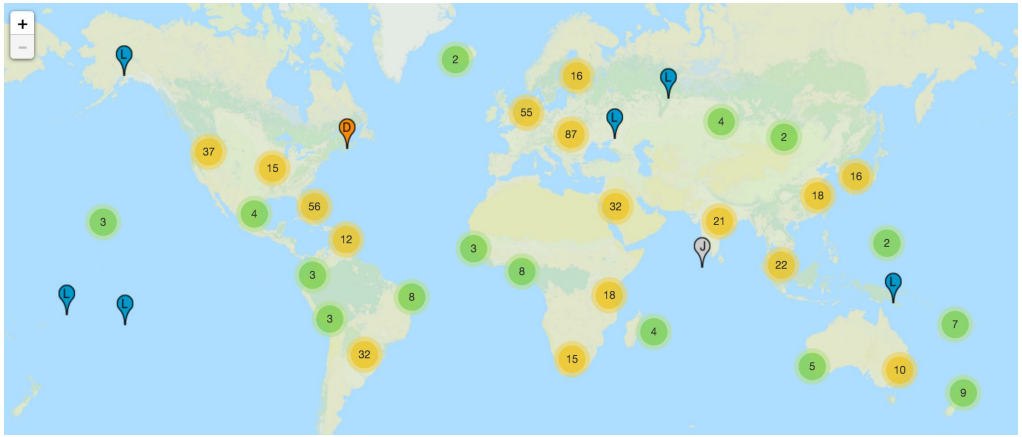

- Authoritative DNS servers for top level “dot” (like com. / net. etc)
- Knows authoritative DNS server of each of TLD & its glue
- Logically 13 servers (from a to m)
- Heavily anycast across hundreds of servers located across the world

What is anycast?

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

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 probes 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. In case of no CHAOS support (in 3/13 roots), it's hard to determine location just based on traceroute due to missing PTRs.
3. Low number of probes in many countries gives less accurate data.
4. Due to blocked ICMP & chaos query - G root was excluded from comparison.
5. Diversification of RIPE Atlas probes across ASNs is not well tested. It is there in some cases while missing in other cases.
6. A considerable number of RIPE Atlas probes are on DSL and may have their 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 in South Asia

1. Afghanistan
2. India
3. Pakistan
4. Nepal
5. Bangladesh
6. Bhutan
7. Sri Lanka
8. Maldives

Countries with root DNS servers

1. ~~Afghanistan~~ *(No active DNS server at the time of testing!)*
2. India (9)
3. Pakistan (4)
4. Nepal (3)
5. Bangladesh (4)
6. Bhutan (1)
7. ~~Sri Lanka~~ *(No active DNS server at the time of testing!)*
8. Maldives (1)

Countries with active RIPE Atlas Probes

1. Afghanistan (0)
2. India (50)
3. Pakistan (7)
4. Nepal (15)
5. Bangladesh (20)
6. Bhutan (11)
7. Sri Lanka (22)
8. Maldives (11)

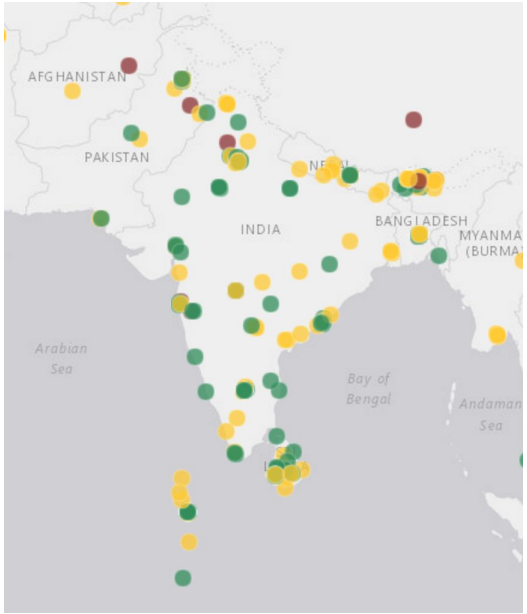

Root DNS Anycast in India

1. J root - Delhi
2. K root - Noida
3. L root - Kolkata
4. J root - Kolkata
5. F root - Chennai
6. D root - Mumbai
7. J root - Mumbai
8. I root - Mumbai
9. L root - Mumbai

Root DNS Anycast in India

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	No	264.33	0.00%	80.43%	20%
B root	No	267.93	0.00%	0.00%	100.00%
C root	No	199.01	0.00%	17%	83%
D root	Yes	123.44	0.00%	80.00%	20.00%
E root	No	137.83	0.00%	78%	21.74%
F root	Yes	60.61	100.00%	0.00%	0.00%
G root	No		0.00%		
H root	No	285.59	0.00%	0.00%	100.00%
I root	Yes	47.97	100.00%	0.00%	0.00%
J root	Yes	73.51	100.00%	0.00%	0.00%
K root	Yes	232.11	0.00%	0.00%	100.00%
L root	Yes	59.96	100.00%	0.00%	0.00%
M root	No	180.03	0.00%	13.64%	86.36%

Root DNS Anycast in India

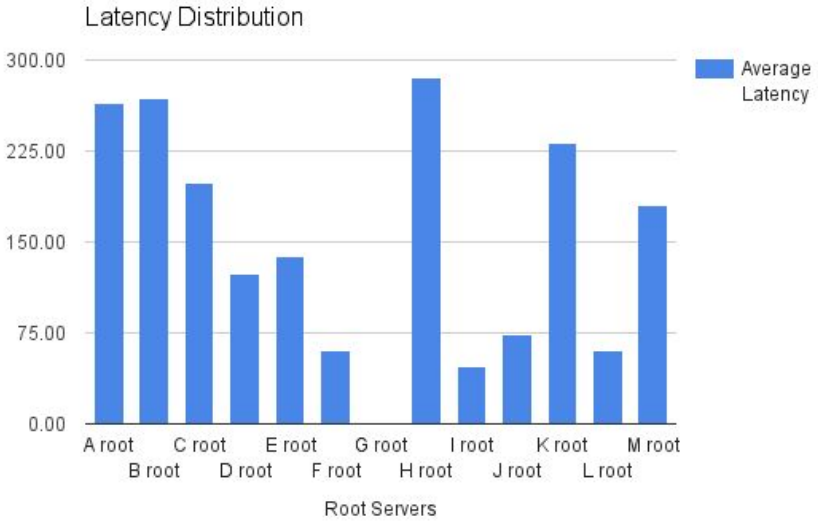

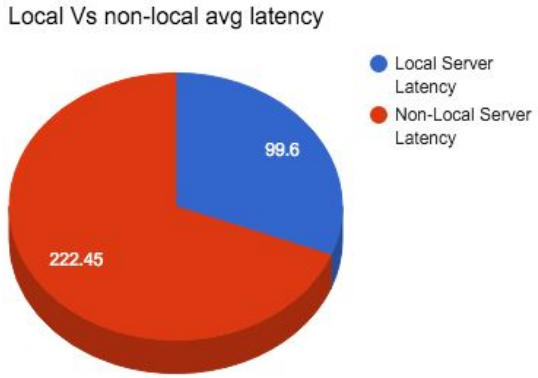

Root DNS Anycast in Pakistan

1. L root - Islamabad
2. L root - Lahore
3. F root - Karachi
4. I root - Karachi

Root DNS Anycast in Pakistan

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	No	196.18	0	0	100%
B root	No	319.26	0	0	100%
C root	No	186.72	0	0	100%
D root	No	246.71	0	0	100%
E root	No	171.66	0	0	100%
F root	Yes	97.84	50%	0	50%
G root	No	n/a	0		
H root	No	265.22	0	0	100%
I root	Yes	117.05	33.3	0	66.6%
J root	No	221.51	0	0	100%
K root	No	208.04	0	0	100%
L root	Yes	24.77	100%	0	0
M root	No	186.26	0	0	100%

Root DNS Anycast in Pakistan

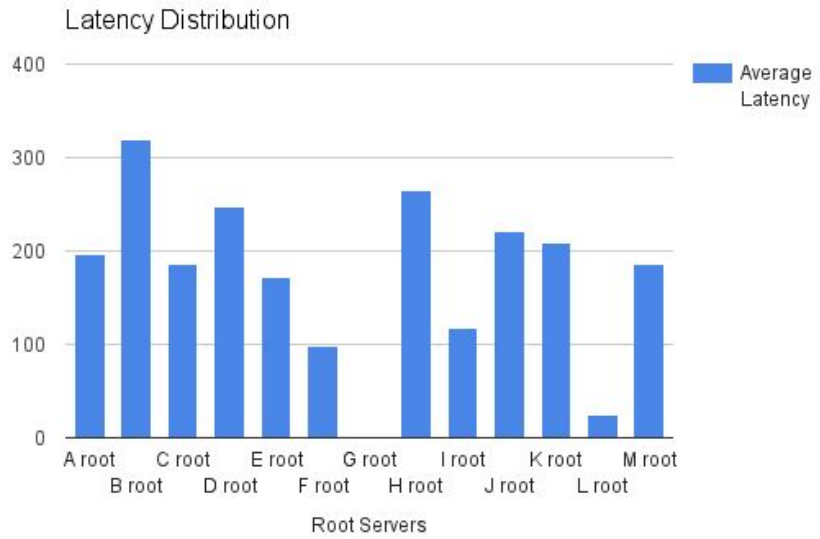

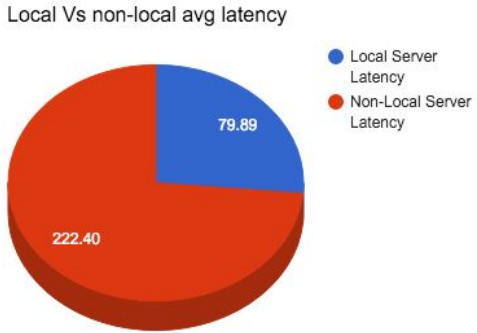

Root DNS Anycast in Nepal

1. D root - Kathmandu
2. E root - Kathmandu
3. I root - Kathmandu

Root DNS Anycast in Nepal

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	No	240.88	0.00%	0.00%	100%
B root	No	274.4	0.00%	0.00%	100%
C root	No	209.88	0.00%	0.00%	100%
D root	Yes	98.14	53.33%	26.67%	20%
E root	Yes	23.07	80.00%	20.00%	0
F root	No	265.04	0%	7.14%	93%
G root	No				
H root	No	333.25	0	0	100%
I root	Yes	89.11	46.67%	46.67%	6.67%
J root	No	217.85	0.00%	0.00%	100%
K root	No	209.91	0.00%	0.00%	100%
L root	No	208.3	0.00%	53.33%	46.67%
M root	No	188.75	0.00%	0.00%	100%

Root DNS Anycast in Nepal

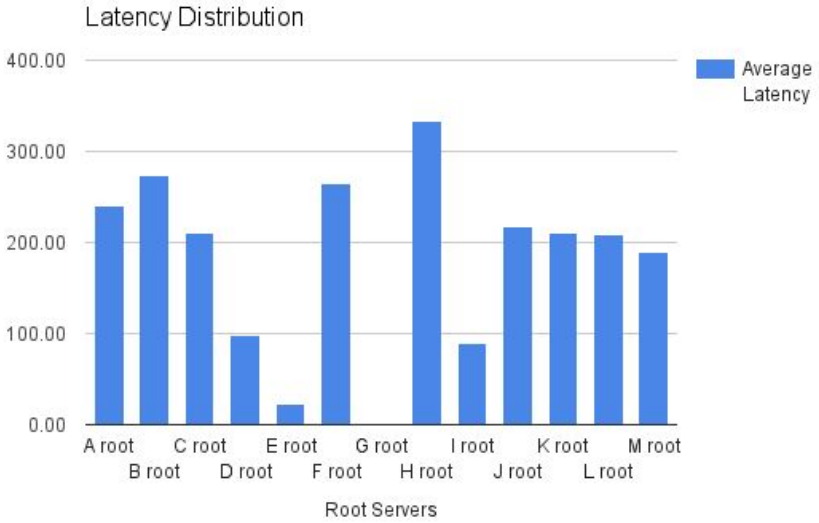

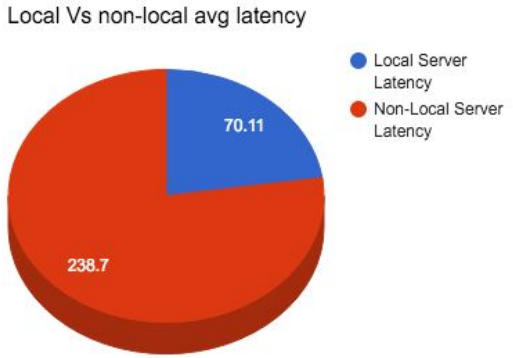

Root DNS Anycast in Bangladesh

1. D root - Dhaka
2. E root - Dhaka
3. F root - Dhaka
4. J root - Dhaka

Root DNS Anycast in Bangladesh

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	No	169.42	0.00%	35.71	74%
B root	No	276.77	0.00%	0.00%	100%
C root	No	209.83	0.00%	0.00%	100%
D root	Yes	75.36	42.11%	47.37%	11%
E root	Yes	70.89	47.06%	52.94%	0.00%
F root	Yes	105.78	53%	0.00%	47%
G root	No				
H root	No	289.95	0.00%	0.00%	100%
I root	No	153.07	0%	47.37%	52.63%
J root	Yes	61.3	50.00%	45.00%	5%
K root	No	196.21	0.00%	0.00%	100%
L root	No	232.24	0.00%	31.58%	68.42%
M root	No	161.29	0.00%	26.32%	74%

Root DNS Anycast in Bangladesh

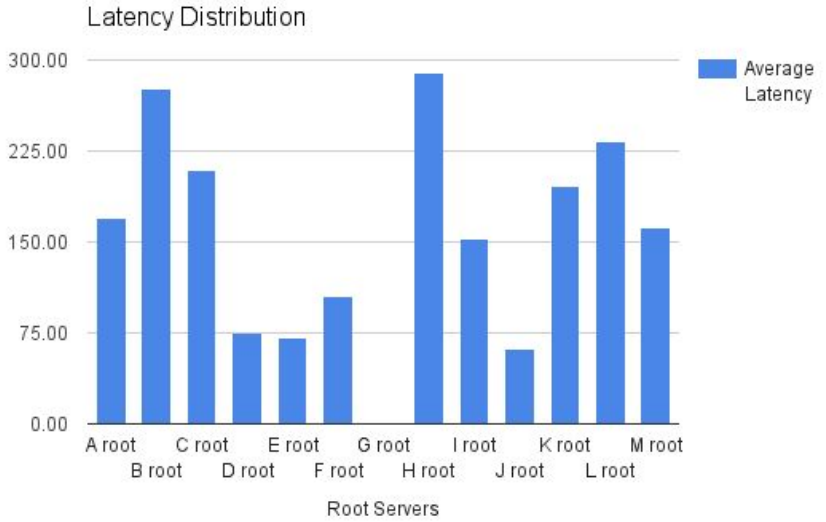

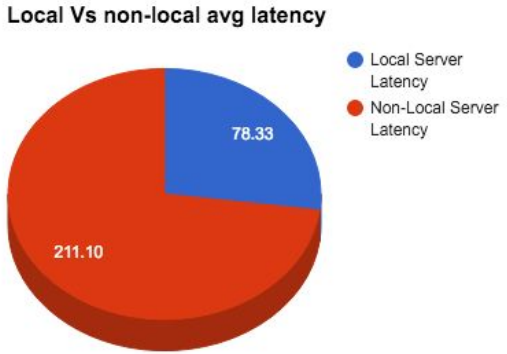

Root DNS Anycast in Bhutan

1. I root - Thimphu

Root DNS Anycast in Bhutan

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	No	151.43	0.00%	78.57%	21%
B root	No	363.65	0.00%	0.00%	100.00%
C root	No	260.53	0.00%	0.00%	100.00%
D root	No	101.65	0.00%	100.00%	0.00%
E root	No	101.29	0.00%	100%	100.00%
F root	No	156.18	0.00%	78.57%	21%
G root	No				
H root	No	320.33	0.00%	0.00%	100.00%
I root	Yes	53.59	100%	0.00%	0.00%
J root	No	167.52	0.00%	45.00%	55%
K root	No	192.85	0.00%	0.00%	100.00%
L root	No	233.47	0.00%	30.00%	70.00%
M root	No	181.77	0.00%	0.00%	100.00%

Root DNS Anycast in Bhutan

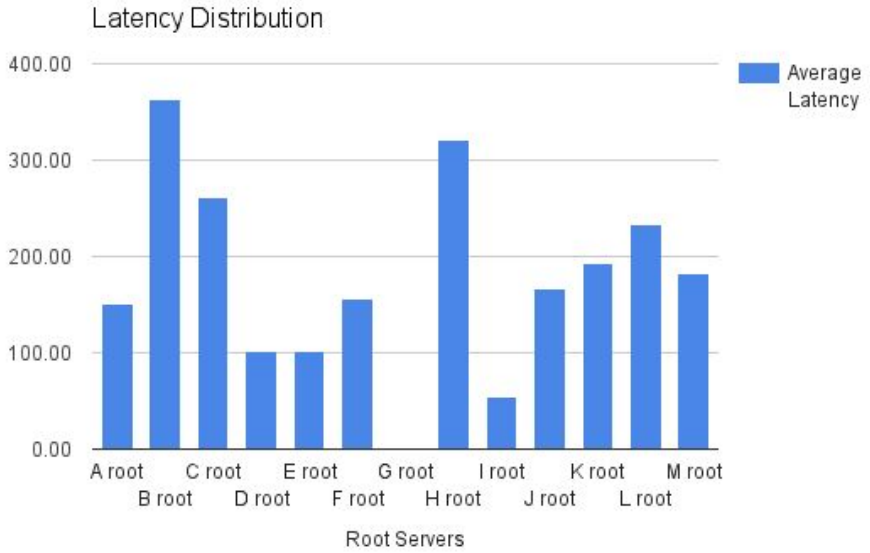

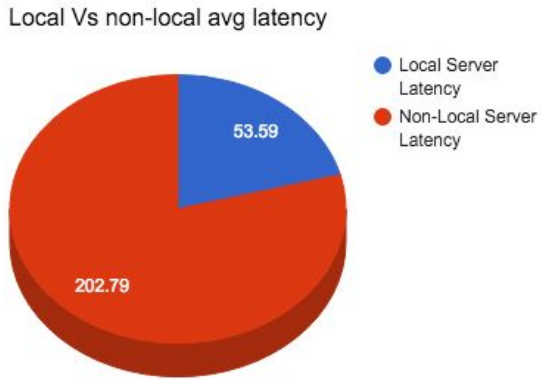

Root DNS Anycast in Sri Lanka

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	No	158.95	0.00%	0.00%	100%
B root	No	272.03	0.00%	0.00%	100%
C root	No	165.78	0.00%	0.00%	100%
D root	No	294.46	0.00%	14.29%	85.71%
E root	No	235.29	0.00%	27%	72.73%
F root	No	192.07	0.00%	0.00%	100%
G root	No		0.00%		
H root	No	312.32	0.00%	0.00%	100%
I root	No	204.84	0.00%	33.33%	66.67%
J root	No	97.64	0.00%	80.95%	19%
K root	No	212.47	0.00%	0.00%	100%
L root	No	165.36	0.00%	4.55%	95.45%
M root	No	163.43	0.00%	13.64%	86.36%

Root DNS Anycast in Sri Lanka

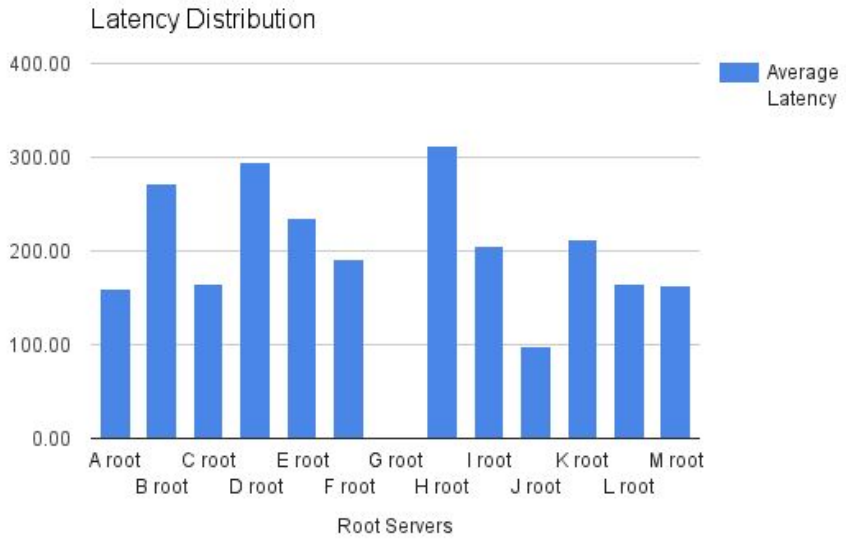

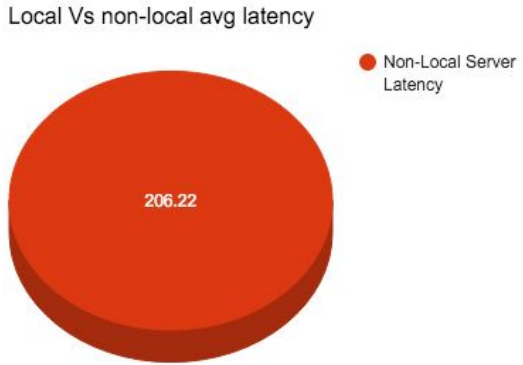

Root DNS Anycast in Maldives

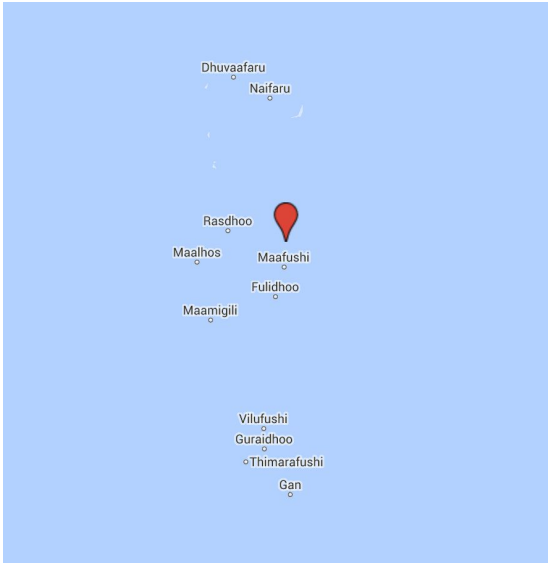

1. J root - Male

Root DNS Anycast in Maldives

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	No	226.01	0.00%	0.00%	100%
B root	No	333.17	0.00%	0.00%	100%
C root	No	259.8	0.00%	0.00%	100%
D root	No	99	0.00%	87.50%	13%
E root	No	90.73	0.00%	100%	0.00%
F root	No	83.53	0.00%	85.71%	14%
G root	No		0.00%		
H root	No	301.13	0.00%	0.00%	100%
I root	No	87.88	0.00%	87.50%	12.50%
J root	Yes	67.22	50.00%	12.50%	38%
K root	No	190.09	0.00%	0.00%	100%
L root	No	265.91	0.00%	0.00%	100%
M root	No	178.5	0.00%	45.00%	56%

Root DNS Anycast in Maldives

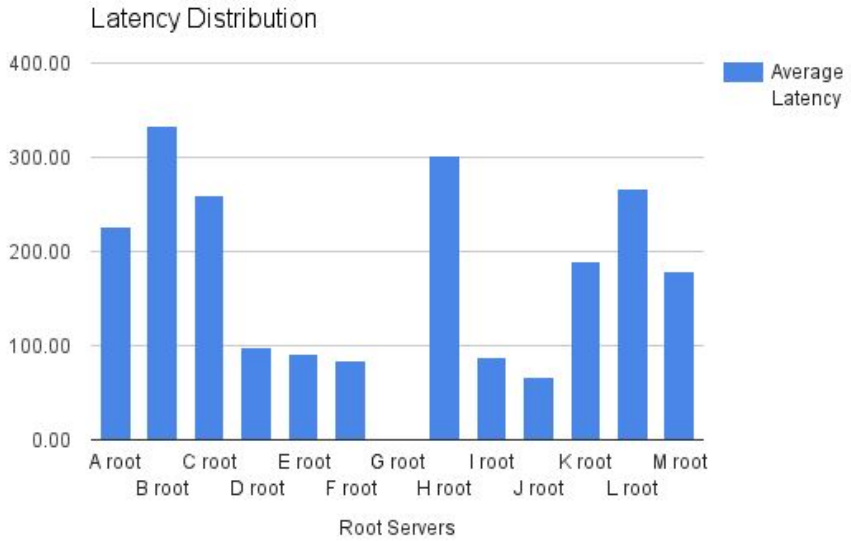

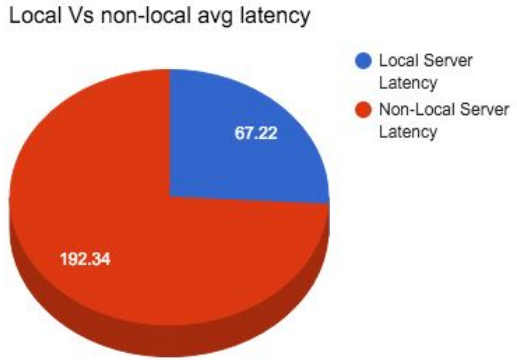

Country wise latency distribution

Country	Average Latency	Number of DNS servers
India	161.03	9
Pakistan	186.77	4
Nepal	196.55	3
Bangladesh	166.84	4
Bhutan	190.36	1
Sri Lanka	206.22	0
Maldives	181.91	1

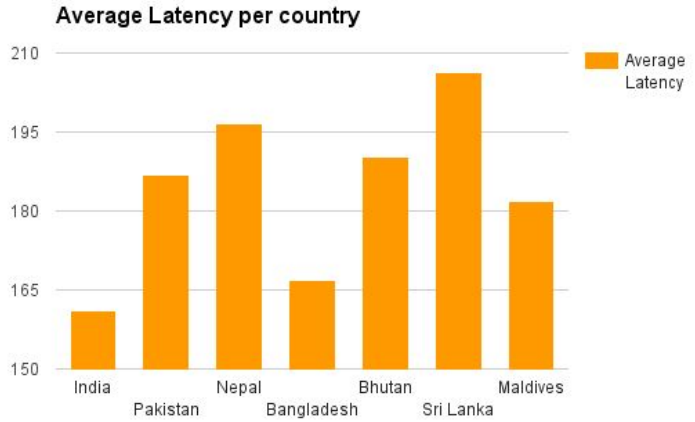

But what about Japan...?

Root DNS Anycast in Japan

1. F root - Osaka
2. M root - Osaka
3. D root - Tokyo
4. G root - Tokyo
5. K root - Tokyo
6. E root - Tokyo
7. I root - Tokyo
8. L root - Tokyo
9. J root - Tokyo
10. M root - Tokyo

Root DNS Anycast in Japan

Root Server	Locally Present	Average Latency	% of DNS Traffic within Country	% of DNS Traffic within region	% of DNS Traffic outside region
A root	No	94.23	0.00%	61.70%	38.30%
B root	No	121.36	0.00%	0.00%	100.00%
C root	No	124.89	0.00%	0.00%	100.00%
D root	Yes	77.55	41.86%	34.88%	23%
E root	Yes	60.63	45.00%	40.00%	15.00%
F root	Yes	51.8	82.93%	4.88%	17.07%
G root	Yes		0.00%	0.00%	100.00%
H root	No	197.41	0.00%	0.00%	100.00%
I root	Yes	16.08	95.74%	4.26%	4.26%
J root	Yes	59.31	41.30%	52.17%	6.52%
K root	Yes	35.07	91.30%	0.00%	8.70%
L root	No	145.53	2.17%	2.17%	96%
M root	Yes	22.1	95.83%	0.00%	4.17%

Root DNS Anycast in Japan

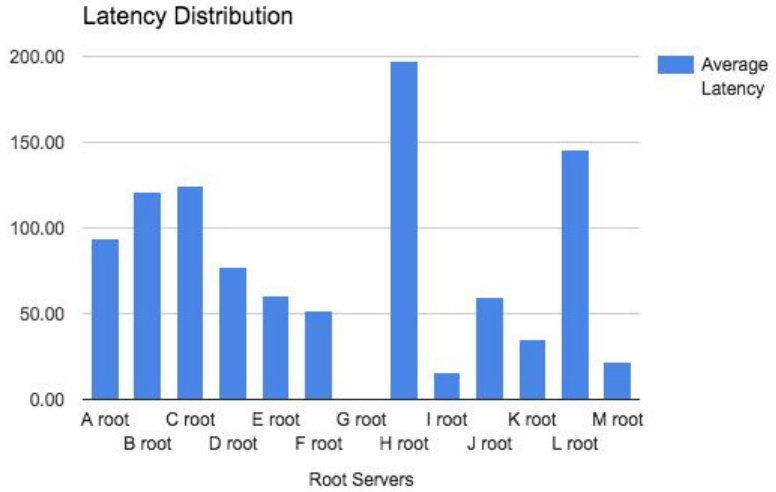

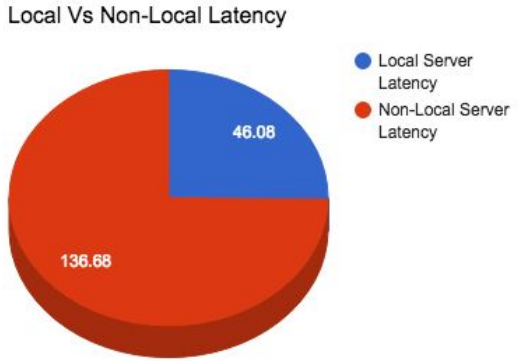

Why anycast fails?

1. Networks often prefer a customer path 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 pick it up based on it’s performance & hence low latency with just one or more root DNS server helps in overall resolution time.

3. Users not using their local ISP's server & relying on popular open DNS recursors have different resolution path alltogether.

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

Conclusions

1. More root DNS servers are good for country. 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.

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

Thankyou!

Questions?

Peering?

Email: anurag@he.net
Twitter: @anurag_bhatia

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