Do You Really Like Me? Anycast Latency and Root DNS Popularity

(abstract, presented at DINR, DNS and Internet Naming Research Directions, 2021)

John Heidemann (1,2) Giovane C. M. Moura (3) Wes Hardaker (1)
1: USC/ISI; 2: USC/Computer Science Dept.; 3: SIDN Labs
1,2: Los Angeles; 3: Amsterdam, 1,2: United States; 3: the Netherlands

It is well known that, when given a choice, many DNS recursive resolvers will favor authoritative servers with lower latency. This performance optimization has been a part of many DNS resolver implementations since the 1990s, and the behavior has been documented in two studies: Yu et al. examined implementations and replayed traces [13], showing that bind prefers lower latencies, although DNSCache, Unbound, and Windows DNS do not. Müller et al studied the Root DNS from thousands of RIPE Atlas nodes [8], finding that 60–70% of traffic to .nl shows at least a weak preference for lower latency. In addition, many DNS services have deployed anycast, in part to reduce latency [3, 4, 10].

The contribution of this abstract to show *how lower DNS latency shifts traffic from one server to another* While prior studies examined DNS from the perspective of a client, we consider the server-side view.

If anycast deployments vary in latency, than implication of a recursive's preference for lower latency is that more traffic will shift to the lower-latency anycast service. We confirm that lower latency results in increased traffic from recursive resolvers that have a choice between multiple anycast service addresses providing the same zone. (This question differs from studies that examine the optimality of a specific anycast service with multiple sites [5–7].)

To examine this question we use public RSSAC-002 statistics for the root server system [9, 12]. From this we use the "traffic-volume" statistic, which reports queries per day for each root anycast service. (Recall that the Root DNS is provided by 13 different anycast service addresses per IP version, each using a different anycast infrastructure.) We show 6 months of data here (2019-11-01 to 2020-05-31), but we noticed similar trends since 2016. This analysis omits G- and I-Root, which did not provide data during this period.

Figure 1 shows the fraction of traffic that goes to each anycast service in the root server system for one year. Two root server identifiers ("letters") deployed additional sites over this period: B-Root originally had 2 sites but added 3 sites in 2020-02-01 [1], then optimized routing around 2020-04-01 [2]. H-Root originally had 2 sites but deployed 4 additional sites on 2020-02-11 and 3 additional sites on 2020-04-06 [11]. While other letters also added sites, B and

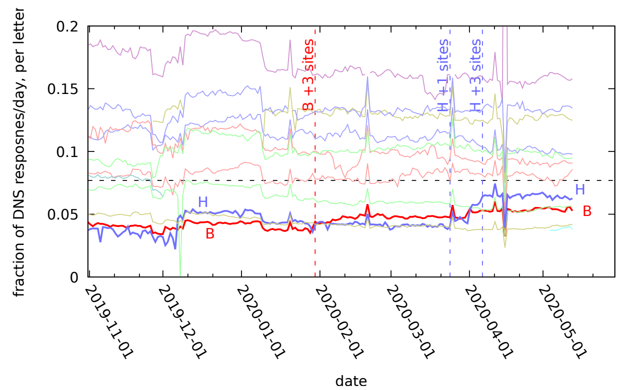

Figure 1: Fraction of traffic going to each root anycast service, per day, from RSSAC-002 data. B- and H-Root are bold lines.

H's changes were the largest improvements relative to their prior size. We see that, B and H's share rises from about 4% in 2019-11 to about 6% in 2020-05.

This data confirms that when new sites deployed by one of the root letters, they offer some clients lower latency for that letter. Lower latency causes some clients to shift more of their traffic to this letter (automatically, as described in [8]), so its share of traffic relative to the others grows.

This data quantifies the long-term uneven balance of traffic across the 13 root letters.

Finally, it suggests that anycast DNS deployments that want to balance traffic across multiple IP anycast deployments (each on its own NS record and IP address) should either keep the size and connectivity of each anycast deployment similar, or expect that load will be uneven.

Acknowledgments: We thank Duane Wessels for making available an aggregated and cleaned collection Root RSSAC data at Github [12].

John Heidemann's research in this paper is supported in part by the DHS HSARPA Cyber Security Division via contract number HSHQDC-17-R-B0004-TTA.02-0006-I (PAADDOS), and NSF CNS-1925737 (DIINER), and by NWO.

Giovane C. M. Moura's research in this paper is supported by the Conconrdia Project, an European Union's Horizon 2020 Research and Innovation program under Grant Agreement No 830927.

REFERENCES

[1] B-Root. B-root begins service from Singapore, Virginia, and Amsterdam. blog post https://b.root-servers.org/news/2020/01/28/new-sites.html, January 2020.

[2] B-Root. B-root optimizes routing. Personal communication from the B-Root operators, 2020.

[3] Matt Calder, Ashley Flavel, Ethan Katz-Bassett, Ratul Mahajan, and Jitendra Padhye. Analyzing the performance of an anycast CDN. In *Proceedings of the ACM Internet Measurement Conference*, pages 531–537, Tokyo, Japan, October 2015. ACM.

[4] Danilo Cicalese and Dario Rossi. A longitudinal study of IP anycast. *ACM Computer Communication Review*, 48(1):10–18, January 2018.

[5] Thomas Koch, Ke Li, Calvin Ardi, Ethan Katz-Bassett, Matt Calder, and John Heidemann. Anycast in context: A tale of two systems. In *Proceedings of the ACM SIGCOMM Conference*, Virtual, August 2021. ACM.

[6] Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. Internet anycast: Performance, problems, and potential. In *Proceedings of the ACM SIGCOMM Conference*, pages 59–73, Budapest, Hungary, August 2018. ACM.

[7] Jinjin Liang, Jian Jiang Haixin Duan Kang Li, and Jianping Wu. Measuring query latency of top level DNS servers. In *Proceedings of the Passive and Active Measurement Conference*, Hong Kong, China, March 2013. Springer.

[8] Moritz Müller, Giovane C. M. Moura, Ricardo de O. Schmidt, and John Heidemann. Recursives in the wild: Engineering authoritative DNS servers. In *Proceedings of the ACM Internet Measurement Conference*, pages 489–495, London, UK, 2017. ACM.

[9] Root Operators. http://www.root-servers.org, April 2016.

[10] Ricardo de O. Schmidt, John Heidemann, and Jan Harm Kuipers. Anycast latency: How many sites are enough? In *Proceedings of the Passive and Active Measurement Conference*, pages 188–200, Sydney, Australia, March 2017. Springer.

[11] The H Root Operators. The U.S. Army Research Laboratory DNS root server. website https://h.root-servers.org, 2020.

[12] Duane Wessles. Rssac002-data. Github respository https://github.com/rssac-caucus/RSSAC002-data/, May 2017.

[13] Yingdi Yu, Duane Wessels, Matt Larson, and Lixia Zhang. Authority server selection of DNS caching resolvers. *ACM Computer Communication Review*, 42(2):80–86, April 2012.

Do You Really Like Me?
Anycast Latency and Root DNS Popularity

John Heidemann[1,2], Giovane C. M. Moura[3], Wes Hardaker[1]

[1]U. of Southern California / Information Sciences Institute and [2]CS Dept. and [3]SIDN Labs

DINR Workshop, 2021-11-16

This research is sponsored by NSF -1925737 (DIINER) and DHS HSHQDC-17-R-B0004-TTA.02-006-I (PAADDOS). The views herein are those of the authors and do not necessarily represent those of DHS or the U.S. Gov't.

DNS recursive resolvers often favor the fastest response

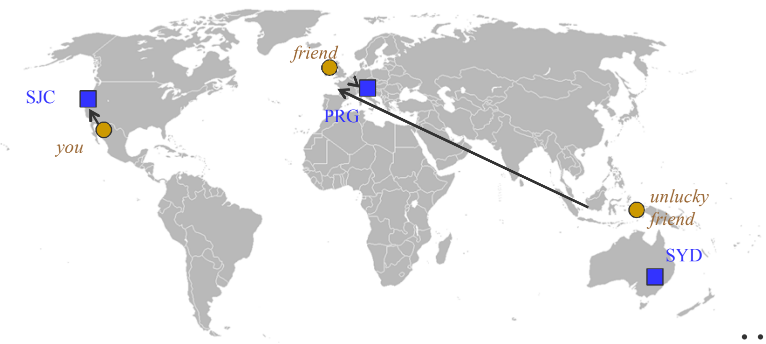

(not all resolver software favors,

and for software that favors, it's not for every query,

and sometimes routing is strange)

…does this matter to authoritative servers?

DNS Recursives Like Fast

- code: Yu et al, "Authority Server Selection", ACM CCR 2012
  - and look at BIND, nsd, knot resolver, source
- experiments from RIPE Atlas: Müller et al., "Recursives in the Wild", IMC 2017

these were from the client side

- even though "fast" doesn't really matter to people for root DNS
  - Koch et al., "Anycast in Context…", SIGCOMM 2021

Us: What About at the Authoritatives?

- if
- "fast" matters,
  we should see unbalanced traffic
  across different NS instances for a given service

what about the authoritative server side?

- can we see differences?
- do they matter?

Methodology: Study Root DNS

- use public RSSAC-002 data
  – daily query counts
- for the root server system
  – 13 different anycast services (RSOs)
  – each with different footprints and changes
- here: 6 months (2019-11 to 2020-05)
- compare to public statements about service change

Query Counts Over Time

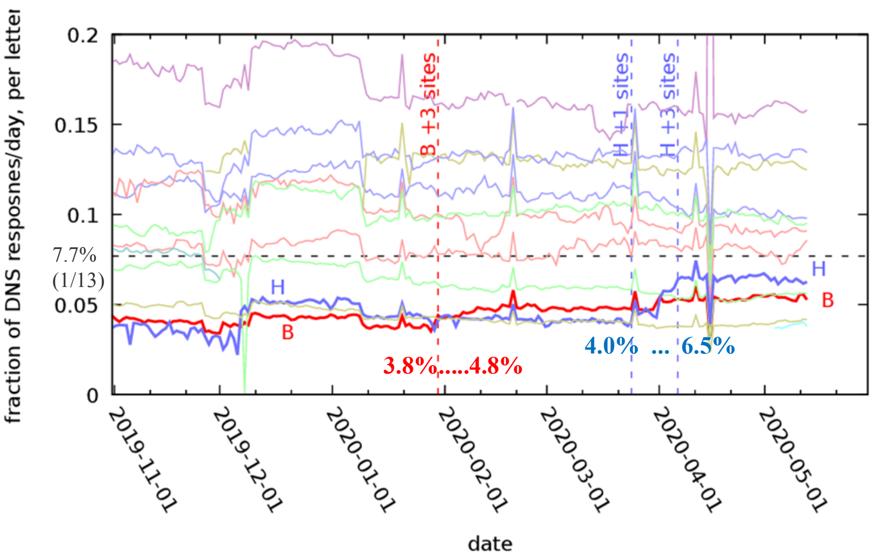

public changes:

B-Root 2020-01-28 + 3 sites
=> +1% of traffic

H-Root 2020-03-22 + 1 site
-0404 + 3 sites
⇒ +2.5% of traffic

(missing data:
G before 2020-05-04,
I after 2019-12-01)

Implications

- changes in traffic may indicate changes in service deployment
  - or the *effectiveness* of new deployments
- root traffic is skewed
  - 1/13th is 7.7%
  - but truth is 3.5% to 18% (1/2 to 2.4x "expected")
  - => RSSAC data can help renormalize
  - (for example, if you're studying DITL and one RSO is missing)

Conclusions

- researchers should consider RSO skew in DITL analysis
- RSO: yes, improvements are appreciated

More Details

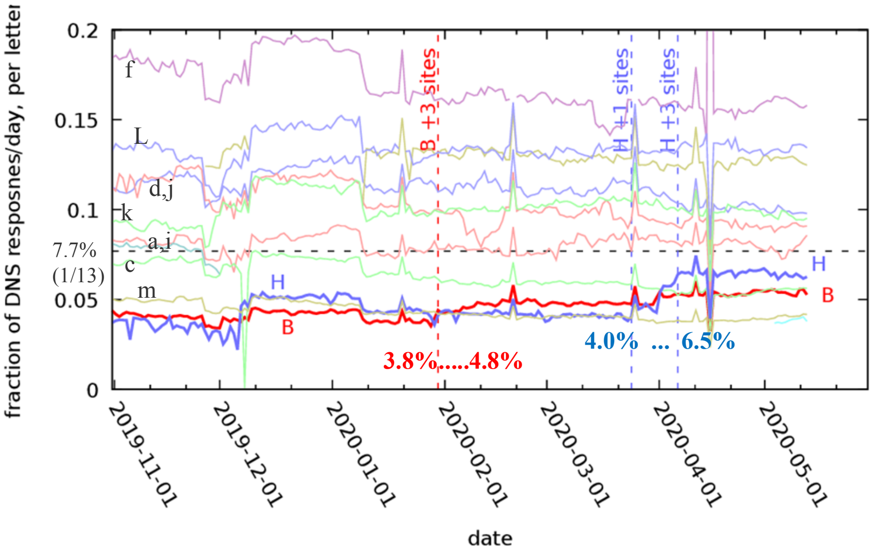

public changes:

B-Root 2020-01-28 + 3 sites
=> +1% of traffic

H-Root 2020-03-22 + 1 site
-0404 + 3 sites
⇒ +2.5% of traffic

(missing data:
G before 2020-05-04,
I after 2019-12-01)