Measuring Query Latency of Top Level DNS Servers

Jinjin Liang[1,2], Jian Jiang[1,2], Haixin Duan[1,2], Kang Li[3], and Jianping Wu[1,2]

[1] Institute for Network Science and Cyberspace, Tsinghua University
[2] Tsinghua National Laboratory on Information Science and Technology
[3] Department of Computer Science, University of Georgia

Abstract. We surveyed the latency of upper DNS hierarchy from 19593 vantage points around the world to investigate the impact of uneven distribution of top level DNS servers on end-user latency. Our findings included: 1) generally top level DNS servers served Internet users efficiently, with median latency 20.26ms for root, 42.64ms for `.com/.net`, 39.07ms for `.org`; 2) quality of service was uneven, Europe and North America were the best while Africa and South America were 3 to 6 times worse; 3) most of the root servers performed well in Europe and North America, but only F, J, L roots showed low query latency in other continents; 4) query latency of F and L roots showed that only about 60% resolvers were routed to the nearest anycast instances. We also revealed two problems that lead to constantly large query latency (6s∼18s) for resolvers. One was buggy implementation of some resolvers on IPv4/IPv6 dual-stack hosts, the other was misconfigured middle-boxes that filtered large or fragmented DNSSEC packets.

1 Introduction

The Domain Name System (DNS) is a fundamental component of Internet which translates domain names into IP addresses for most of the Internet applications. This makes DNS query latency a critical factor that affects the quality of Internet experienced by end-users.

Essentially, DNS is a distributed database organized as a hierarchical tree. On the top of the hierarchy are root zone and some top level domain (TLD) zones such as `.com` and `.net`. DNS authority servers of top level zones are crucial for Internet operation since these servers serve Internet users all over the world as the start points of the whole domain name space. Their performance is also important to end-users as some popular implementations could still visit top level DNS servers frequently even with the local caches [10].

A common technique to increase DNS robustness and performance is DNS zone replication, in which one DNS zone can be served by multiple authority servers in various locations. For example, root zone is currently served by 13 logical root servers and hundreds of anycast instances.

An important issue of deploying replications of top level zones is geographic distribution. Historically, most of the root servers were located in the

M. Roughan and R. Chang (Eds.) PAM 2013, LNCS 7799, pp. 145–154, 2013.

United States. Internet users of other regions inevitably stood longer DNS query latency due to geographic distance. Previous research indicated that Europe and Asia were underprovisioned while North America was overprovisioned [7], suggesting that some root servers should be relocated.

Recently, the deployment of replication instances for top level zones grows massively with wide adoption of anycast, which further allows multiple instances in different locations to use a same IP address. Publicly available information [1] shows 319 anycast instances for 13 root servers have been deployed all over the world. Our curiosity is that how Internet users from different regions experience this progress and whether the uneven quality of service has been improved. One recent article compared the number of root instances and the population served in different continents and concluded that the distribution of root servers was still very uneven [1]. However, the conclusion of [1] is based on simplistic statistics of users served per root server. We hope to have a more technical and comprehensive evaluation of the quality of services of upper DNS servers.

We probed 19,593 open recursive resolvers to query top level servers using a method we called *NXDOMAIN-Query* and King technique [5]. *NXDOMAIN-Query* could obtain the overall latency from resolvers to root or TLD level while King technique could measure the latency between a resolver and an arbitrary nameserver. With the measured latency, we compared the DNS performance in different regions and further analyzed the current state of top level zones' replication deployment. We find: 1) generally top level DNS servers serve Internet users efficiently, with median latency 20.26ms for root, 42.64ms for `.com`/`.net`(they share the same infrastructure), 39.07ms for `.org`; 2) quality of service is still uneven: Europe and North America are the best while Africa and South America are 3 to 6 times worse; 3) in Europe and North America, most of the root servers perform well, but in other continents only F, J, L roots show low query latency; 4) query latencies of F and L roots show that only about 60% resolvers are routed to the nearest anycast root instances.

Along with the results for initial motivation, we also observed anomalous large latency from a group of resolvers, ranging from 6s to 18s. Our further investigation revealed two reasons. One was buggy implementation of some resolvers on IPv4/IPv6 dual-stack hosts. The other was misconfigured middle-boxes on certain paths which filtered large or fragmented DNSSEC responses.

2 Methodology

We collect plenty of open recursive resolvers and then drive these resolvers to query our targets. Using the round trip times (RTTs) observed from these resolvers, we further estimate DNS query latency between these resolvers and the targets. This approach has two advantages: 1) we do not need the direct control of our vantage points, which allows us to scale our study extensively; 2) probing open resolvers triggers *real* DNS lookup behaviors in the wild, which helps us to observe anomalous behaviors and further identify the causes.

---

[1] `http://www.root-servers.org`

2.1 Collecting Open Recursive Resolvers

We collect 19593 open resolvers in total by three ways: 1) extracting open resolvers from the query log of a busy DNS authority nameserver (42%); 2) probing DNS authority servers of Alexa top 1M sites (42%); 3) inquiring help from other researchers (16%). [2] It is worth to note that we must exclude DNS forwarders since they would query their upper resolvers rather than query authority servers directly. The geographic distribution of these resolvers is detailed in Table 1.

Table 1. Distribution of Open Resolvers (Based on GeoIP database)

Continent	# of countries	# of ASes	# of resolvers	% of total
Europe	45	2821	7169	36.59
North America	25	1837	5525	28.20
Asia	40	940	6056	30.91
South America	11	173	426	2.17
Oceania	7	131	248	1.27
Africa	26	77	149	0.76
Unknown	-	-	20	0.10
Total	154	5979	19593	100.00

2.2 NXDOMAIN-Query Technique

We utilize a method called *NXDOMAIN-Query* to indirectly measure the query latency from a resolver to a domain level. The main idea of NXDOMAIN-Query is leveraging non-existent domain names to control recursive resolvers to stop at specified domain levels. Besides, *NXDOMAIN-Query* uses a fresh non-existent domain name for each request to avoid cached negative responses. For example, to measure the query latency from a resolver to root level, we issue the resolver with a DNS query containing a fresh non-existent TLD from a client. When the resolver receives this query, it asks one of the root servers, then receives a *NXDOMAIN* response and replies to the client with the *NXDOMAIN* answer. Assume T_{c-root} is the whole query latency observed from the client, T_{r-root} is the latency between the resolver and the root server, T_{c-r} is the latency between the client and the resolver which can be measured by issuing the resolver with a non-recursive query. Then we can estimate T_{r-root} by $T_{r-root} = T_{c-root} - T_{c-r}$. Similarly, we can also measure DNS lookup latency from resolvers to TLD level.

The limitation of *NXDOMAIN-Query* is that this method only allows us to measure the overall latency from selected resolver to a group of DNS servers serving a specified level of DNS domain, rather than to a certain DNS server. The reason is that different resolvers implement different server selection strategies which cannot be controlled indirectly.

---

[2] The limitation of using open resolvers as vantage points is that query latency may be affected by network condition of the resolvers. However, our exploration verifies that about 70% of these resolvers are also authority servers, thus we believe networks may only have limited influence on our measurement results.

2.3 King Technique

We leverage King technique [5] to indirectly measure DNS lookup latency from selected resolver to one certain DNS server, as a complement to the above *NXDOMAIN-Query*. The basic idea of King technique is tricking recursive resolver to query any designated IP address through pointing the nameserver of a controllable domain to that IP address, and then estimating latency between the resolver and the designated IP with the observed RTTs. Please refer to the original paper [5] for more technical details.

3 Measurements and Results

Using methods introduced above, we investigated the impact of geographic distribution of top level DNS servers by measuring overall query latency of both root and TLD level, as well as the individual query latency of each of the 13 root servers from different regions. We also analyzed anycast proximity of F and L roots and compared the differences among continents [3].

3.1 Query Latency of Root and TLD Hierarchy

We assessed the overall query latency of root level and popular TLDs through *NXDOMAIN-Query* approach. For each open resolver we collected, we measured its query latency to root, `.com`/`.net`, `.org` respectively. To reduce measurement noises, we continuously measured the latency over 500 times during a two days period for each resolver and used median value as a resolver's final latency.

Figure 1 shows the CDF of all resolvers' query latency to the three measured top DNS zones. We can see that generally these three zones serve global users efficiently. For most of the resolvers the latency are small, and the median latency for all three zones are less than 50ms. Specifically, root zone has the lowest overall query latency and the median latency is about 20.26ms. `.org` slightly outperforms `.com`/`.net`, the median latency for `.org` is 39.07ms while for `.com`/`.net` is 42.64ms.

A surprising result from Figure 1 is that a few resolvers constantly show very large query latency, with the values mostly around 6, 8, 12, 18 seconds. We present our investigation for this strange behavior in Section 4.

We break down the measurement results to compare the differences among various continents. Figure 2 shows the quartiles of all resolvers' latency to the three zones in each continent. The results show that the cost of querying the three top DNS zones is uneven across continents. All six continents can be categorized into three groups. Most of the resolvers in Europe (EU) and North America (NA) have distinctly smaller latency than other continents, especially comparing to resolvers in South America (SA) and Africa (AF), the median latency of which are 3 to 6 times larger. Asia (AS) and Oceania (OC) are more complicated. While the median latency can be equal (OC) or slightly larger than those of

---

[3] The results are available at `https://github.com/dnsmeasurement/latency`.

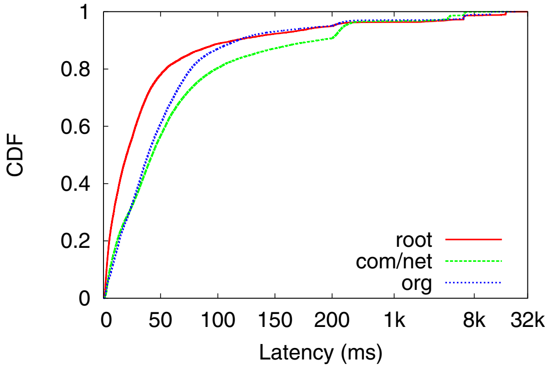

Fig. 1. Cumulative distribution of latency of root, `.com/.net`, `.org`

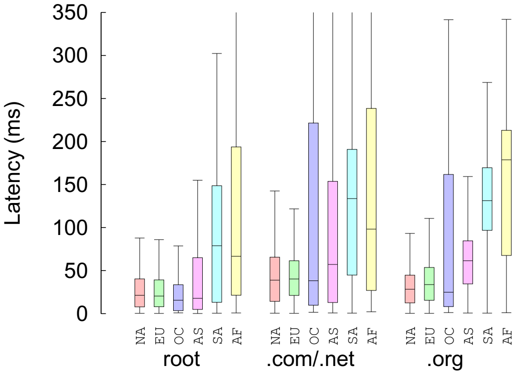

Fig. 2. DNS query latency of root, `.com/.net`, `.org`, breaking down by continent

EU and NA, the quartile values are usually much bigger. This indicates the quality of service of top level DNS servers varies greatly among countries or autonomous systems (ASes). We leave the detailed analysis of country-level or AS-level differences as our future work.

3.2 Query Latency of Thirteen Root Servers

Using King technique, we measured the query latency from each open resolver to each of the thirteen root servers. Same as above, we launched the measurement over 300 times in two days continuously and for each root server we extracted the median value as each resolver's final query latency.

We categorize the resolvers by continents and show all their query latency to the thirteen root servers in Figure 3. Since the number of resolvers in each continent is different, we normalize the results in Figure 3 before analyzing them.

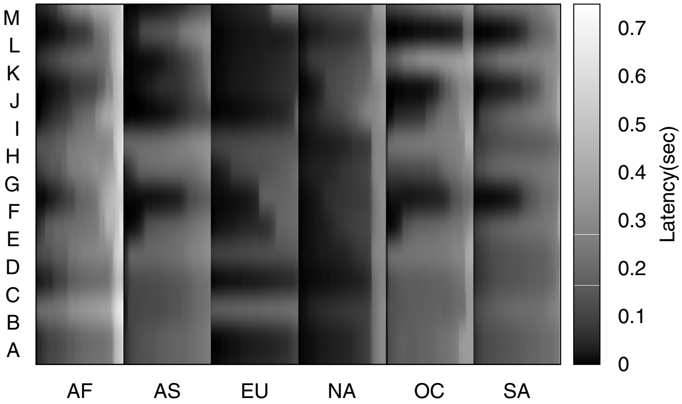

Fig. 3. Query latency of 13 root servers in different continents

In Figure 3, the most remarkable root servers are F, J, L. They perform well in all of the continents. Their median query latency for each continent are all below 200ms. On the contrary, B root performs the worst among all the roots. Most of its query latency are over 300ms in all continents except North America. The results are consistent with public information of root server deployment. F, J, L root anycast nodes have been deployed widely all over the world while B root only has one nodes in America.

Figure 3 also shows the uneven performance of root servers in different continents. Europe and North America are the best and their query latency to most root servers are quite small. By contrast, latency of most roots in Africa, Oceania and South America are much larger. Especially in Africa, query latency of some roots, like B and G, are even over 600ms. This result reflects the current state of root anycast deployment that Europe and North America deploy much more root instances than other continents do.

Compared with the result in Section 3.1, we observe that although not all the roots perform well in each continents, the overall latency of root level are relatively small. This results from the resolvers' server selection mechanisms which usually choose a best server to request.

3.3 Proximity of Root Anycast

Anycast instances of a logical root server share a same anycast address, but have different unicast addresses for management. Comparing query latency of anycast address with the minimum latency for all the unicast addresses could infer whether a resolver is indeed routed to the nearest anycast node. In previous research, this property was referred to as *anycast proximity* [2].

We used King technique to measure the anycast proximity of F and L roots, whose unicast addresses were publicly available. Our measurement repeated about 200 times in two days and used medians as the final latency for each

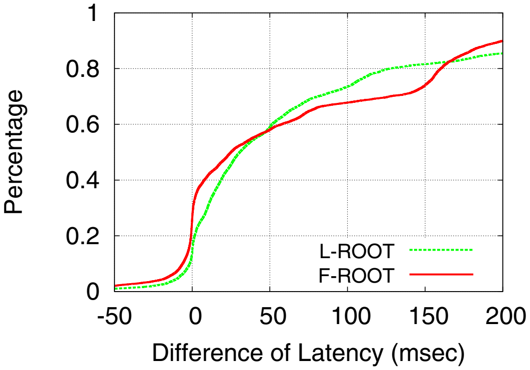

(a) Overall proximity

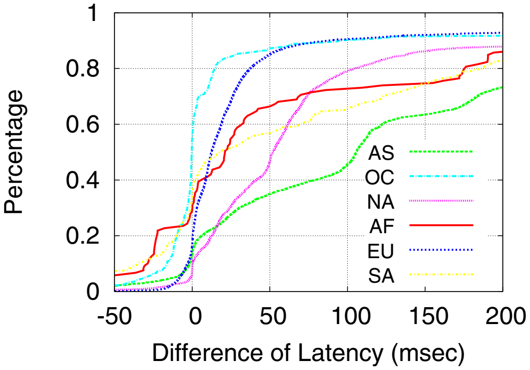

(b) L root proximity

Fig. 4. The proximity of root anycast

resolver. Finally, we computed the proximity of a root $T_{proximity}$ by $T_{anycast} - min(T_{unicast})$, where $T_{anycast}$ was the latency of a root's anycast address and $min(T_{unicast})$ was the minimum latency of all that root's unicast addresses.

Figure 4(a) shows the overall proximity of F and L roots. We find that a fair fraction of resolvers are not routed to the anycast nodes closest to them. For example, about 40% of the resolvers are routed to servers more than 50ms farther away from the nearest anycast nodes for both F and L roots. This is most likely caused by routing policies like BGP and the hierarchical deployment [4]. What's more, we observe that about 2% and 1% of resolvers whose $T_{proximity}$ are below -30ms for F root and L root, which means that the queries are routed to servers that are nearer than the closest nodes. Several possible reasons could lead to such strange phenomenon: errors in measurement results, paths for anycast are

---

[4] F-root are deployed hierarchically, only 2 of 49 nodes advertise the anycast prefix globally.

faster than that for unicast to a server or missing some unicast nodes in our experiment list (e.g. lack of timely update or masquerading roots [4]). We leave this investigation as our future work.

We also classified the resolvers by continents to analyze the anycast proximity in different regions. Figure 4(b) shows the proximity of L roots in six continents. We see that the quality of anycast proximity in Oceania and Europe outperforms those in other continents for L root, only 13% and 15% of the resolvers are directed to the nodes 50ms farther than the closest servers. On the other hand, resolvers in Asia suffer the worst quality of proximity, $T_{proximity}$ for 65% of its resolvers are over 50ms.

4 The Cause of Large Query Latency

In Section 3.1, we observe that a group of open resolvers (totally 664, 3.2% of all) constantly show very large query latency (larger than 2 seconds) when visiting root and TLDs. For root, the latency are mainly around 6, 18 seconds, while for TLDs, the latency are 4, 6 seconds and 6, 12 seconds for `.com/.net` and `.org` respectively.

After exploring these problematic resolvers, we find out two causes that are responsible for the large latency: buggy implementation of certain resolvers on IPv4/IPv6 dual-stack hosts and misconfigured middle-boxes on certain paths which filter large or fragmented DNSSEC responses.

4.1 Buggy Implementation on IPv4/IPv6 Dual-Stack

We first focus on resolvers that consume 18 seconds constantly when traversing root level. We use *fpdns* tool to gather information of these resolvers and find that nearly all of them are running on BIND 9.2. To observe these resolvers' resolution process, we set up a testing domain with three name servers and drive the problematic resolvers to visit our name servers through querying them for subdomains under the testing domain.

We find that everytime we put a new IPv6 address into the glue record, two seconds extra delay will be added. We infer that the large latency are related to the IPv6 address of name servers. 9 of the 13 root servers are configured with IPv6 addresses, so these resolvers need about 18 seconds to traverse the root level. Similarly, 2 `.com/.net` and 6 `.org` TLD servers use IPv6 addresses, which lead to 4 and 12 seconds latency respectively. Further investigation on source code confirms that BIND 9.2.x running on IPv4/IPv6 dual-stack host always prefers IPv6 authorities even if they are unreachable. New versions of BIND ($\geqq$ 9.3) have fixed this problem.

4.2 Filtering of DNSSEC Response

Excluding large latency explained above, the rest ones are mostly around 6 seconds. Using fpdns tool again, we find that most of the rest problematic resolvers are BIND 9.3.x. We also drive them to query our name servers to observe their behaviors.

We notice that when we configure our testing domain with DNSSEC, the 6 seconds' resolution will occur. Observing from our name server, we find that resolvers firstly send 3 queries with `EDNS0` every 2 seconds sequentially and then send a query without `EDNS0` at last. Since DNSSEC is enabled, responses for the first three `EDNS0` queries contain DNSSEC records which are larger than 512 bytes. We infer that large DNSSEC responses are dropped on the paths to these problematic resolvers, which eventually causes 6 seconds resolution latency.

5 Related Work

Quite a few measurements were carried out to study the DNS infrastructure. Brownlee *et al.* [3] and Lee *et al.* [6] investigated the performance indicators of root and some TLD nameservers, such as DNS response time and request loss rate. Liston *et al.* [8] and Lee *et al.* [7] analyzed the performance impact of top level nameserver placement. Our work measured the query latency to top level DNS servers from a large number of vantage points and tried to correlate the query latency with the current nameserver deployment. Yu *et al.* [11] measured the placement of top level DNS servers. While they focused on identifying the locations of nameservers and assessing their robustness, our work aimed at revealing the impact of nameserver deployment on performance. Sarat *et al.* [9] and Ballani *et al.* [2] measured the availability and the proximity of DNS root anycast, and also provided suggestions for deployment strategies. Our work focused on the proximity of F and L roots anycast and investigated the differences among various regions.

6 Conclusion

DNS is a public resource shared by Internet users all over the world. However, historically, top level DNS servers are unevenly deployed, which leads to unfair quality of DNS service in different regions. Recently top level DNS servers, especially root servers have been deployed massively with wide adoption of anycast. Our measurement shows that this progress improves the overall DNS performance. However the quality is still uneven among different regions. Nevertheless, the adoption of anycast enables rapid deployment of replication in underprovisioned areas. ISPs should be more proactive to deploy local root anycast instances to improve their DNS query performance.

Our measurement also observed anomalous large latency. While the cause of buggy implementation might not be an issue, the other cause of filtering large DNSSEC response is more important. With DNSSEC being a crucial protocol of Internet in future, large DNS response and IP fragment should be considered as regular rather than harmful traffic. The community should take more efforts to measure unexpected DNS packet filtering and discuss possible implications.

Acknowledgements. We are grateful to the anonymous reviewers and our shepherd Simon Leinen for their valuable comments. This work is supported by the National Basic Research Program of China (973 Project, Grant No. 2009CB320505) and National Natural Science Foundation of China (Grant No. 61161140454). Kang Li's research on this work is partially supported by US National Science Foundation (NSF) Office of Cyberinfrastructure grant 1127195.

References

1. The (very) uneven distribution of DNS root servers on the internet, `http://royal.pingdom.com/2012/05/07/the-very-uneven-distribution-of-dns-root-servers-on-the-internet/`
2. Ballani, H., Francis, P., Ratnasamy, S.: A measurement based deployment proposal for ip anycast. In: Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement, IMC 2006, pp. 231–244. ACM, New York (2006)
3. Brownlee, N., Claffy, K., Nemeth, E.: DNS root/gtld performance measurements. In: USENIX LISA, San Diego, CA (2001)
4. Fan, X., Heidemann, J., Govindan, R.: Identifying and characterizing anycast in the domain name system. Tech. rep. (2011)
5. Gummadi, K.P., Saroiu, S., Gribble, S.D.: King: estimating latency between arbitrary internet end hosts. In: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement, IMW 2002, pp. 5–18. ACM, New York (2002)
6. Lee, B.S., Tan, Y.S., Sekiya, Y., Narishige, A., Date, S.: Availability and effectiveness of root DNS servers: A long term study. In: 2010 IEEE Network Operations and Management Symposium (NOMS), pp. 862–865 (April 2010)
7. Lee, T., Huffaker, B., Fomenkov, M., et al.: On the problem of optimization of DNS root servers' placement (2003)
8. Liston, R., Srinivasan, S., Zegura, E.: Diversity in DNS performance measures. In: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurment, IMW 2002, pp. 19–31. ACM, New York (2002)
9. Sarat, S., Pappas, V., Terzis, A.: On the use of anycast in DNS. In: Proceedings.15th International Conference on Computer Communications and Networks, ICCCN 2006, pp. 71–78 (October 2006)
10. Wessels, D., Fomenkov, M., Brownlee, N., Claffy, K.: Measurements and Laboratory Simulations of the Upper DNS Hierarchy. In: Barakat, C., Pratt, I. (eds.) PAM 2004. LNCS, vol. 3015, pp. 147–157. Springer, Heidelberg (2004)
11. Yu, Y., Cai, J., Osterweil, E., Zhang, L.: Measuring the placement of DNS servers in top-level-domain