On the Performance of DNS Resolvers in the IPv6 and ECS Era
(Abstract)

Rami Al-Dalky and Michael Rabinovich
Case Western Reserve University

We revisit the issue of the performance of DNS resolution services available to Internet users. While several prior studies [1,6,7] addressed this important issue, significant developments, namely, the IPv6 finally getting traction [5] and the adoption of the ECS extension to DNS [4] by major DNS resolution services, warrant a reassessment under these new realities. DNS interactions presage user communication with content servers, and thus DNS latency in answering queries can affect responsiveness of user communication (even though there is a growing evidence that DNS prefetching is effective at masking this latency). Moreover, content delivery networks commonly use DNS to route user communication to nearby edge servers. Given the prevalence of CDNs in the modern Web[1], the ability of the DNS system to effectively facilitate user mapping has a profound effect on the end-user experience. Consequently, we focus on both these aspects of DNS behavior: its latency in answering queries and the quality of user-to-edge-server mappings. Further, with major CDNs and DNS resolvers now supporting IPv6 and most client devices being dual-stacked, we consider how the choice of IP version affects those behavioral aspects.

To conduct this study, we use 200 dual-stack RIPE Atlas probes from the ~1600 probes listed as dual-stack by RIPE Atlas, which we chose based on their stability and diversity of represented autonomous systems and geographic locations. Using more probes would not appreciably improve general representativeness of our results because of a general strong skew of RIPE Atlas probes towards North American and, especially, European locations, which are already disproportionally represented in our sample. Excluding probes that did not consistently produce measurements, we ended up with 188 productive probes distributed in 74 countries across 188 ASes and 6 regions: 52 in North America (US and Canada), 70 in Europe, 38 in Asia, 11 in Latin America, 7 in Africa, and 10 in Oceania (a region that includes Australia and Pacific islands). Further, 36 of these probes with high probability used ISP-provided resolvers as their default resolvers (which we deduced from the fact that the probe and its resolver either shared the same /24 address block or come from the same autonomous system according to Team Cymru). These 36 probes represented 23 countries in all 6 regions, and 36 autonomous systems. We use these probes to compare the performance of public and ISP resolvers. Using our 188 vantage points, we measure DNS latency of resolving hostnames from several DNS resolution services and TCP latency of client-to-edge-server mappings of several CDNs. We attempt to focus on client-to-resolver latency by pre-warming the resolver cache before taking the measurement. Since this turned out to not always ensure a cached response, we also separate our measurements into hits and misses based on the TTL of the DNS records received. To drive our measurements, we use the first 100 websites from Majestic top-1M list[2] that (a) support both IPv4 and IPv6 protocols, (b) are accelerated by a CDN -- which we determine by examining the CNAME chain of the DNS resolution of a website with "www" prepended, and (c) support HTTPS such that we can measure the latency between the probes and the assigned edge server[3]. With these websites, we are able to measure Akamai (65 sites), Cloudfront (17), Google (9), Fastly (8), and Incapsula (1). The DNS resolutions services we assess include Google Public DNS, OpenDNS, Cloudflare, and Quad9. Our key findings include the following:

- We find that DNS resolution services differ drastically -- by an order of magnitude in some locations -- in their query response time. In particular, we find established resolvers (Google DNS and OpenDNS) to lag far behind relative newcomers (Cloudflare and Quad9) in terms of DNS latency, and trace the cause to drastically lower cache hit rates, which we further trace to less cache sharing within the resolver platform. Curiously, we also find that, when it comes to resolving Akamai-accelerated names, Cloudflare *misses* are often faster than Google's and OpenDNS's *hits*. We attribute this to a combination of wide footprints of both Cloudflare resolvers and Akamai's authoritative DNS platforms, which may result in a shorter aggregate network path of a Cloudflare's miss than Google's hit.
- We find that the DNS latency is generally little affected by whether the client requests hostname resolution to an IPv4 or IPv6 address, and by the client's choice of IP version to interact with the DNS system.
- While multiple prior studies found that public resolvers produce worse CDN user mappings than ISP resolvers [1,6,7], we find that public resolvers have largely closed the gap with ISP resolvers in this regard.
- Finally, in most locations, we observe IPv6 penalty in the latency of client-to-CDN-edge-server mappings produced by the resolvers. Moreover, this penalty, while often substantial, still does not rise above typical thresholds (250msec as recommended default) employed by the Happy Eyeballs algorithm [8] for switching to IPv4 communication. Thus, dual-stacked clients in these locations may stay in IPv6 and experience suboptimal performance. This finding complements and contrasts with a prior study [3] that probed Alexa top-10K websites from 80 vantage points and found insignificant latency difference, with 91% of these sites having IPv6 latencies within 1msec of their IPv4 counterparts.

Further details of this study can be found in [2] and our measurement data is available at [9]. This work is supported by NSF through grant CNS-1647145.

---

[1] According to Cisco estimates, over half of all web traffic is delivered through a CDN. See https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-indexvni/complete-white-paper-c11-481360.html# Toc484813991

[2] http://majestic.com/reports/majestic-million

[3] RIPE Atlas does not support HTTP requests but does allow a TLS handshake, thus enabling the TCP latency measurement.

References

[1] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. Comparing DNS Resolvers in the Wild. In ACM IMC, 2010.

[2] R. Al-Dalky and M. Rabinovich. Revisiting Comparative Performance of DNS Resolvers in the IPv6 and ECS Era. http://arxiv.org/abs/2007.00651

[3] V. Bajpai and J. Schönwälder. Measuring the Effects of Happy Eyeballs. In ACM Applied Networking Research Workshop, 2016.

[4] C. Contavalli, W. van der Gaast, D. Lawrence, and W. Kumari. Client Subnet in DNS Queries. RFC 7871, May 2016.

[5] J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, and M. Bailey. Measuring IPv6 Adoption. In ACM SIGCOMM, 2014.

[6] H. Hours, E. Biersack, P. Loiseau, A. Finamore, and M. Mellia. A study of the impact of DNS resolvers on CDN performance using a causal approach. Computer Networks, 109, 2016.

[7] C. Huang, D. A. Maltz, J. Li, and A. Greenberg. Public DNS system and Global Traffic Management. In INFOCOM, 2011.

[8] D. Schinazi and T. Pauly. Happy Eyeballs Version 2: Better Connectivity Using Concurrency. RFC 8305, December 2017.

[9] https://www.dropbox.com/sh/ioswha9fha3gmxn/AABVNt78urpVshVBbQvAGfvKa?dl=0