End-User Mapping: Next Generation Request Routing for Content Delivery

Nitish Paradkar, Xinghao Li

EECS 589 Paper Review

1 Full Reference

Chen, F., Sitaraman, R.K., and Torres, M., "End-user mapping: Next generation request routing for content delivery," *Proc. of ACM SIGCOMM '15*, 45(4):167-181, Oct. 2015.

2 Paper Summary

In this paper, the authors describe how they leveraged the EDNS0 extension to learn about the first 24 bits of the client IP address and redirect them to a closer edge data center. The main motivation behind this is that there are several cases where clients are not close to their LDNS (most notably when using a public resolver), and the traditional NS-based mapping results in many clients being directed to edge servers that are far away, resulting in unnecessary network overhead. The authors begin by finding out how close a client geographically tends to be to a LDNS on average. They find that this number varies greatly (on the order of 5000 miles), and that clients who use public resolvers as their LDNS tend to have significantly higher client-LDNS distances. They then analyze how well the rollout of their new client-based mapping worked for end users and found that there were significant decreases in mapping distance (8 times) and RTT times (2 times) for clients who were far away from their LDNS, as well as more modest reductions in mapping distance and RTT times for clients that were already close to LDNS. The authors also conduct some research into how getting fewer than 24 bits of the client IP address affects client cluster sizes. They discover that getting more bits of the client IP address leads to a higher number of client clusters and a smaller cluster geographical radius. They also finally conclude that this new client based mapping system will help more for CDNs that are deployed in more locations than CDNs deployed in fewer locations.

3 Highlights

One of the things that we like about this paper is how they used user-data to find out how far away clients tend to be from their LDNS. The benefits to doing this are twofold. On the one hand, their results show that there are a significant number of users who are far away geographically from their LDNS, and this helps give reason to why some other method for deciding how to map a user to an edge server besides just using the traditional NS-based mapping is important. It also shows the reader where a rollout of end user mapping would benefit users the most (such as users in Vietnam, India and Turkey in Figure 6 of the paper). On a similar note, we also appreciate how the authors showed detailed data on how far client-LDNS distances tend to be for clients using public resolvers (Figure 7), as it again gives a good idea on how much farther public resolvers tend to be on average from their clients. It also helps show which types of clients would benefit more from the user-based mapping. While there have been previous efforts to show how far away clients tend to be from their LDNS in a network sense [3], we ultimately also appreciate the geographic distance metric as it helps us solidify these distances in our mind.

We also liked that the authors used metrics like time to first byte (TTFB) and content download time in their evaluation. While they probably don't directly measure the effectiveness of the proposed mapping technique (as the geographical distance and RTT cover this well), we feel like these are ultimately very important numbers as they determine how much this new technique actually helps the end user. It is often the case that edge servers have to fetch data from the origin before forwarding it to the client (especially in the case of dynamic websites). So, while this metric may not directly measure how effective the user based mapping technique is, it ultimately shows the expected performance gain that an end user could expect when downloading a real web page.

We also like how the authors were able to take advantage of the EDNS0 protocol to gain more information about the client. The EDNS0 protocol was proposed in 1999, and has been used fairly widely as it is required for DNSSEC [2]. In addition to this, both OpenDNS and Google DNS [4], two of the biggest public resolvers, have already implemented this protocol. This is extremely useful as it does not require each LDNS to implement a new protocol. Thus providing the benefit of making this mapping technique very feasible to use.

We also appreciate how the authors include some details into how selecting a set of IP blocks affects the client cluster sizes and radii. Using intuition, it is easy to deduce that as more coarse IP blocks are selected, the number of clusters decreases and the cluster radius increases; however, it is still nice to see some numerical data to see numerical data showing this pattern. Including this information is important as there are several cases where one LDNS could give each clients the address to different edge servers. If this happens, it needs to store each unique entry into its cache. As more coarse IP blocks are used, it intuitively makes sense that less cache entries would be used per query. By including how many clusters are formed based on the prefix length selected, the reader can start to get some sense of the tradeoff between selecting a large IP prefix length, and having too many entries in the cache for one query.

Finally, we also appreciate that the authors deployed this mapping technique and tested out how well this works for a wide variety of users in their performance analysis. In many research papers, new protocols tend to be tested in simulations, which give limited feedback into how useful these methods can be. By deploying the protocol in production and measuring the results, the authors make a very strong case for why end user mapping is a useful extension to their existing mapping technique.

4 Possible Improvements and Extensions

In general, this paper does a great job on exploring the performance of the end-user mapping by analysing the difference between traditional NS-based mapping and the new end-user mapping with a series of thorough and insightful measurements. However, in our opinion, there are some points that the paper does not thoroughly consider.

First of all, the paper does not mention how well EDNS0 is deployed. In order for end-user mapping to work well, it is necessary to have sufficient LDNS servers that are fully compatible with EDNS0. If only a small group of LDNS servers support EDNS0, the new mapping method will not end up providing much of a benefit (even if the performance is exceptional). Based on some quick research, somewhere between 60-85% of LDNS servers fully support EDNS [1]. This does not consider how well LDNS in very widespread AS's support EDNS0. This would be an example of clients that would benefit greatly from end-user mapping. Thus, we feel like the authors could have made a more complete the case for end-user mapping by including some data on EDNS0 adoption rates .

Another thing to note is that in this paper, the authors only consider geographical distance between clients and their LDNS. While geographical distance generally correlates positively with network distance, it is also possible that a client and LDNS that are close to one another could have a very long circuitous network path. We were confused as to why the authors did not also include some network related metric, such as RTT, in their analysis of client-LDNS distances. There have been some studies done that show that only 64% of clients are in the same AS as their LDNS, and only 16% of these client-LDNS pairs are in the same network-aware cluster [3]. We feel like the inclusion of this type of network related information such as this could have been incredibly useful as it could have helped paint a more complete picture of client LDNS distances.

The authors show some analysis for client clusters as a function of the IP prefix blocks. In general, we see that as the number of bits used in the client prefix increases, the client cluster radius decreases and the number of client clusters increases. While this information is useful, it is important to note that when using end-user mapping, each LDNS now caches the DNS query for both the query name and the client

IP prefix. Intuitively, as the number of bits in the client prefix used in EDNS0 increases, there are more likely to be cache misses for a LDNS. On every cache miss, the LDNS will have to do additional work to resolve the client query by doing the full DNS traversal to find the right IP address to correctly map the client to an edge server. In addition to this, it is even possible that two clients with slightly different IP prefixes will map to the same edge server (but the LDNS will now have two cache entries for this query). Thus, there is a tradeoff in choosing the correct number of prefix bits to use in end-user mapping. On the one hand, choosing too few bits will result in poor performance for the client, but choosing too many bits in the prefix could introduce additional work for each LDNS. In the paper, the authors already show how DNS queries recorded by Akamai's name servers increased after the rollout of end-user mapping. We feel like it would also be nice to include some details on how changing the number of prefix bits used in the EDNS0 protocol changes the increase in queries recorded by the Akamai edge servers. In addition to this, it would also be useful to see how changing the number of bits in the prefix would also affect the performance of end-user mapping (by again comparing the four metrics of distance, RTT, TTFB, and client download time). We feel like including this information would be useful as it would help find the perfect balance in the trade off between doing unnecessary work for each LDNS and providing optimal performance for end users.

In section 3, the authors did a widespread experiment on collecting the data for the client-LDNS distance in 25 countries around the world. However, one important thing they ignored is that they did not take China (CN) into account. According to the statistics of Internet users on Google Public data[5], there are more than 700 million Internet users in China. This number is definitely the highest over the world, and is possibly more than the sum of the internet users in 25 countries the paper selected. Thus, not including the samples in China may cause the analysis insufficient or unpursuable.

In addition, in section 4, the paper does not explain explicitly about the dataset for evaluating the performance of end-user mapping. I got no information about the distribution for the number of samples in each country. According to the interpretation of CDN in the background section of the paper, we deduce that it providers more benefit in countries or areas with bad outbound internet connection to oversea servers, but less benefit for most developed countries with high total interconnection bandwidths across multiple networks. For example, according to Xinghao's Internet using experience in China, the page loading time for oversea websites are pretty long, and sometimes the entire website is unable to open. One of the most crucial reasons why this occurs is because of the bad global interconnection in China. There are about 20 Tier 1 ISPs around the world, but none of them is in China[6]. Therefore, the oversea data traffic must detour a long way to the corresponding Tier 1 ISP if corresponding peer relationships are not established at Tier 2 level, even if the geographic distance between client and server is much shorter than the distance the data travels in the wire. In this situation, the CDN will provide noticeable benefits with its higher bandwidth volume and priority than its clients. However, in developed countries such as United States and members of European Union, the bandwidth is pretty abundant for most users to get sufficient bandwidth and unnoticeable delay for global web services. Thus, CDN's will provide more limited benefits for these countries. In conclusion, one possible improvement we came up with for this paper is to consider more samples in developing countries with poor oversea Internet connections, especially China. In addition, when discussing the performance boost for end-user mapping in section 4, we think it is better to divide the samples into two sections: developed countries with good Internet quality and developing countries with relative poor Internet services.

In this paper, the authors only analyze data across the Akamai CDN. Nevertheless, there are a number of its competitors such as Cloudflare[7], Google Cloud CDN[8], and Amazon Cloudfront[9], and most of its competitors have also accepted the EDNS0 information from LDNS and did the similar thing as the end-user mapping of Akamai. Thus one possible extension for this paper is adding discussions over other CDN brands besides Akamai. But it may be a little bit hard to access the statistics since the data of some commercial CDNs are considered as secret, and not likely possible to be released out.

5 References

[1] https://www.isc.org/blogs/partial-edns-compliance-hampers-deployment-of-new-dns-features/

[2] https://en.wikipedia.org/wiki/Extension_mechanisms_for_DNS

[3] Z. Morley Mao, C. Cranor, F. Douglis, M. Rabinovich, O. Spatscheck, and J. Wang. A precise and efficient evaluation of the proximity between web clients and their local DNS name servers. In Proceedings of the Usenix Annual Technical Conference, 2002, Monterey, CA, June 2002

[4] http://www.cdnplanet.com/blog/which-cdns-support-edns-client-subnet/#cdns

[5] https://www.google.com/publicdata/explore?ds=d5bncppjof8f9_&met_y=it_net_user_p2&idim=country:CHN:IND:USA&hl=en&dl=en

[6] https://en.wikipedia.org/wiki/Tier_1_network

[7] https://www.cloudflare.com/

[8] https://cloud.google.com/cdn/

[9] https://aws.amazon.com/cn/cloudfront/