

Finding the best answer: measuring the optimization of public and authoritative DNS

Jia ZHANG^{1*}, Haixin DUAN¹, Jian JIANG², Jinjin LIANG³ & Jianping WU^{1,4}

¹*Institute for Network Science and Cyberspace, Tsinghua University, Beijing 100084, China;*

²*International Computer Science Institute, UC Berkeley, Berkeley 94720, USA;*

³*Qihoo 360 Technology Co. Ltd, Beijing 100015, China;*

⁴*Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China*

Received 13 October 2017/Revised 22 December 2017/Accepted 26 February 2018/Published online 1 November 2018

Citation Zhang J, Duan H X, Jiang J, et al. Finding the best answer: measuring the optimization of public and authoritative DNS. *Sci China Inf Sci*, 2019, 62(3): 039107, <https://doi.org/10.1007/s11432-017-9370-2>

Dear editor,

Public domain name system (DNS) and authoritative DNS service providers have adopted various methods to optimize the accuracy of their resolution [1]. The public DNS server supports EDNS client subnet (ECS) [2], which can send the subnet information of the end users to the authoritative DNS server, and the authoritative DNS also gradually supports ECS to schedule IPs based on the ECS data information. In addition, some CDN service providers use anycast IP to schedule resources and reduce dependence on the authoritative DNS IP schedule [3].

Owing to technical, business, and legal considerations, only a part of these optimizations are in use. Previous researchers have studied the advantages of a single optimization technology; however, the end users of the Internet need to understand the overall optimization effect, existing research on which is not very clear.

The aim of this study is to measure the deployment of different optimization techniques in the authoritative DNS and CDN servers. Compared with the initially proposed technologies, deployment of these current technologies has substantially increased in the Alexa Top 10K sites [4].

ECS measurement. Since the ECS deployment in public DNS servers is widely known and does not need special measurement techniques, the ECS

deployment in authoritative DNS servers was principally measured.

The `dig` measurement mainly involved measuring the consistency of the resolution results to find whether the authoritative DNS servers support ECS. The procedure is as follows.

(1) Find more than 1 million servers from around the world that support public DNS resolution via network scanning and select about 20000 servers based on the scan results. Servers for each B network segment(/16) located in 5 continents and about 190 countries were selected.

(2) Use `dig` to get the authoritative DNS servers of the Alexa Top 10K domains. If the NS record of a domain maps to multiple IPs, only the first IP address of the authoritative DNS server is chosen.

(3) Use `dig` to directly get the resolution results of the Alexa Top 10K domains from their authoritative DNS servers using the "`client`" option; the value of "`client`" is the IP address of the public DNS server selected in Step (1) (`netmask=24`).

(4) Use `dig` to get the resolution results of the Alexa Top 10K domains from their public DNS servers selected in Step (1).

For one domain, if the two conditions are met, that is, if the value of the "`client`" option obtained from Step (3) is equal to the DNS IP obtained from Step (4), the resolution result is the same in Steps (3) and (4); whereas if the value of

* Corresponding author (email: zhangjia@cernet.edu.cn)

the "client" option in Step (3) is unequal to the DNS IP in Step (4), the resolution result of the two steps is different; thus, we can conclude that the authoritative DNS server supports ECS.

In the actual measurement, even if the DNS server supports ECS, the resolution result may be different in Steps (3) and (4). One domain name can have more than one authoritative DNS server; if the zone file is not synchronized, the resolution consistency will get affected if we have to use different authoritative DNS servers in different regions in our measurement. Additionally, for load balance, some authoritative DNS servers of CDNs may return different results for a request from the same client. For this condition, a threshold of consistent ratio is set. We found that based on the consistent ratio, domains can be classified into four categories:

(1) Consistent ratio $< 10\%$. These domains should deploy CDN but ECS should not be supported.

(2) $10\% < \text{consistent ratio} < 80\%$. In these domains, anycast CDNs or traditional domain names map to limited multiple IPs. The consistency of such domain name resolution results is inversely proportional to the number of mapped IP addresses.

(3) $80\% < \text{consistent ratio} \leq 100\%$. In these domains, the names deployed on CDNs support ECS. Most of the results are consistent because of ECS, and the consistent ratio is higher than the domain names mentioned above; however, due to reasons such as multiple authoritative servers and load balancing, some inconsistent results may exist.

(4) Consistent ratio = 1. These domains are the anycast CDNs or traditional domain names that map to exactly one IP.

In this study, we set a threshold that if more than 80% of the results are the same, and the resolution results with different "client" options are different, the domain will support ECS.

Reflection measurement. Theoretically, dig with "client" options can measure the ECS deployment of each authoritative DNS server. However, in real networks, even if the service providers support ECS in both the public DNS and the authoritative DNS of CDN, its use has some restrictions because of considerations of privacy and commercial interests. For example, Akamai supports only the ECS information forwarded from Google and Open DNS; Google prohibits the forwarding of user-defined ECS information to some authoritative DNS and fills the ECS information by itself. For the DNS servers that do not support dig measurement, we adopt a more effective measurement

method that is inspired by the DNS reflection attack [5].

Suppose clients A and B are two measurement nodes in different regions, and resolve servers A and B are two different resolver nodes of one public DNS supporting ECS (e.g., Google). Then, the resolve server A is close to client A and the resolve server B is close to client B . In this case, the measurement procedure is as follows:

(i) Client A makes a standard DNS request of one domain to a public DNS which supports ECS and gets the resolution result A_1 .

(ii) Client A makes a DNS request of the same domain in Step (i) with the source IP of client B to the same public DNS of Step (i), and the resolution result A_2 is returned to client B .

(iii) Client B makes a standard DNS request of the same domain to the same public DNS of Step (i) and gets the resolution result A_3 .

If the authoritative DNS server supports ECS, it will return the resolution result based on the "subnet"; so $A_1! = A_2$ and $A_2 = A_3$. If the authoritative DNS server does not support ECS, it will return the resolution result based on the unicast IP of the resolve server node. As the unicast IPs are the same in Steps (i) and (ii) but different from the unicast IP in Step (iii) due to the anycast technology in the public DNS, the resolution results are $A_1 = A_2$ but $A_2! = A_3$.

Considering that different public and authoritative DNS servers have different characteristics that support ECS, we use three public DNS servers that support ECS, namely Google, Open DNS, and DNS Pod, to obtain this measurement. If a resolution demonstrates the characteristic of " $A_1! = A_2$ and $A_2 = A_3$ ", we deduce that the authoritative DNS server supports ECS.

Compared with the traditional dig measurement method, the deployment of this new method may be more complicated and time consuming. First, the new method would need networks of measurement nodes to enable the sending DNS query request packets with fake source addresses; Second, the distance between measurement nodes, Clients A and B , would have to be relatively large; for example, one in Asia and the other in America. In addition, the measurement target should be derived from well-known CDN service providers from around the world. If it is a regional CDN, Clients A and B will get the same result in Steps (i) and (iii). On account of these preconditions, there are many restrictions on the selection of measurement nodes and the measurement target. Compared with the traditional measurement method, this method finds fewer DNS servers that support ECS, but can find some DNS servers that do not

support dig measurements. Therefore, although this method cannot measure ECS based on the DNS independently, it is an effective complement to the dig measurement method.

Anycast CDN measurement. Anycast technology in CDN is an optimization method to reduce the dependency on DNS servers and the user location of CDN scheduling. To measure the deployment of anycast technology in CDN, we first use two measuring nodes to directly send DNS requests to the authoritative DNS domain. If the resolution results returned to the two measurement nodes are the same, we use IP `traceroute` to observe the path information, which includes the AS information about each hop from the measurement node to the destination. For the two paths from two different measurement nodes, if the previous AS of the destination is not the same, the two measurement nodes access two different destination nodes with the same IP address. This implies that the corresponding nodes support anycast.

On account of a failure of CDN nodes or routing jitter [6], all requests of different measurement nodes may route to the same destination. To reduce this bias, each domain and IP are measured at least 3 times over different time periods. From a statistical point of view, this specific case does not affect our judgment.

Measurement results. We used the above methods to measure the deployment of each optimization technology in the Alexa Top 10K domains

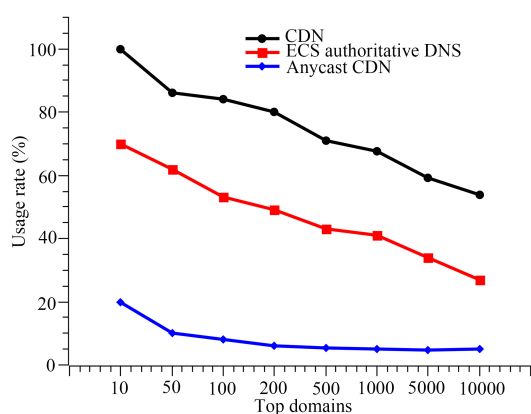


Figure 1 (Color online) Optimization technologies used by DNS and CDNs (*X* axis is not uniform).

from June 2017 to September 2017. To reduce the impact of DNS manipulation on our measurement, we ignored the measurement results of 153 sensitive domain names based on Pearce's work [7].

In Figure 1, the ratio of the Top 10K domains which support CDN, ECS, or anycast IP CDN technology is shown. The CDN usage ratio has not significantly increased since the last five years [8,9]. However, the proportion of authoritative DNS supporting ECS has significantly improved. 70% of the top 10 sites and more than 50% of the top 100 sites support ECS. In addition, about 5% of the Alexa Top 10K sites, 20% of the top 10 sites, and 10% of the top 100 sites support anycast CDN. Therefore, it was concluded that the increasing use of these optimization technologies can effectively increase the CDN resource scheduling accuracy.

Acknowledgements This work was supported by National Key Research and Development Program (Grant No. 2017YFB0803202), and National Natural Science Foundation of China (Grant No. 61472215).

References

- 1 Qin Z, Xiao C, Wang Q, et al. A CDN-based domain name system. *Comput Commun*, 2014, 45: 11–20
- 2 Contavalli C, van der Gaast W, Lawrence D, et al. Client subnet in DNS requests (RFC7871). 2016. <https://tools.ietf.org/html/rfc7871>
- 3 Calder M, Flavel A, Katz-Bassett E, et al. Analyzing the performance of an anycast CDN. In: *Proceedings of the 2015 Internet Measurement Conference (IMC'15)*, New York, 2015. 531–537
- 4 Alexa. The top sites on the web. 2017. <http://www.alexa.com/topsites>
- 5 Anagnostopoulos M, Kambourakis G, Kopanos P, et al. DNS amplification attack revisited. *Comput Secur*, 2013, 39: 475–485
- 6 Xu M W, Li Q, Yang Y, et al. Self-healing routing: failure, modeling and analysis. *Sci China Inf Sci*, 2011, 54: 609–622
- 7 Pearce P, Jones B, Li F, et al. Global measurement of DNS manipulation. In: *Proceedings of the 26th USENIX Security Symposium (USENIX Security 17)*, Vancouver, 2017. 307–323
- 8 Otto J S, Rula J P. Content delivery and the natural evolution of DNS: remote dns trends, performance issues and alternative solutions. In: *Proceedings of the 2012 Internet Measurement Conference (IMC'12)*, Boston, 2012. 523–536
- 9 Streibelt F, Böttger J, Chatzis N, et al. Exploring edns-client-subnet adopters in your free time. In: *Proceedings of the 2013 Internet Measurement Conference (IMC'13)*, New York, 2013. 305–312