`dnstracker`: Measuring Centralization of DNS Infrastructure in the Wild

Luciano Zembruzki[1], Arthur Selle Jacobs[1], Gustavo Spier Landtreter[1],
Lisandro Zambenedetti Granville[1], and Giovane C. M. Moura[2]

[1] Institute of Informatics, Federal University of Rio Grande do Sul
Av. Bento Gonçalves, 9500, Porto Alegre, Brazil
{lzembruzki, asjacobs, gslandtreter, granville}@inf.ufrgs.br
[2] SIDN Labs and TU Delft, the Netherlands
giovane.moura@sidn.nl

Abstract. The Internet Domain Naming System (DNS) is one of the pillars for the Internet and has been the subject of various Distributed Denial-of-Service (DDoS) attacks over the years. As a countermeasure, the DNS infrastructure has been engineered with a series of replication measures, such as relying on multiple authoritative name servers and using IP anycast. Even though these measures have been in place, we have seen that, when servers rely on third-party DNS providers for reliable services, there may be certain levels of infrastructure centralization. In this case, an attack against a DNS target might affect other authoritative DNS servers sharing part of the infrastructure with the intended victim. However, measuring such levels of infrastructure sharing is a daunting task, given that researchers typically do not have access to DNS provider internals. In this paper, we introduce a methodology and associated tool `dnstracker` that allows measuring, to various degrees, the level of both concentration and shared infrastructure using active DNS measurements. As a case study, we analyze the authoritative name servers of all domains of the Alexa Top 1 Million most visited websites. Our results show that, in some cases, up to 12.000 authoritative name servers share the same underlying infrastructure of a third-party DNS provider. As such, in the event of an attack, those authoritative DNS servers have increased the probability of suffering from collateral damage.

Keywords: Domain Name System, Measurements, Centralization

1 Introduction

The Internet Domain Naming System (DNS) provides a globally hierarchical naming space on the Internet that enables the mapping of hosts, networks, and services to IP addresses [9]. As such, DNS is one of the core services of the Internet. To resolve a domain name (*e.g.*, `ufrgs.br`), first, a client sends a DNS query to its *DNS recursive resolver* (resolver hereafter), which is a DNS server that, on behalf of the client, can resolve the domain name. If the resolver does not have a DNS record in a cache, it will query the DNS hierarchy for a response.

Resolvers are responsible for sending queries to *DNS authoritative nameservers* (authoritative DNS server hereafter), which are the servers responsible for providing answers to resolvers about the fetched domain. These authoritative DNS servers are divided into zones and only know the content of a DNS zone from local knowledge, and thus can answer queries about those zones [5].

The authoritative DNS servers have been frequent victims of Distributed Denial-of-Service (DDoS) attacks. The Root Zone, which is authoritative for the Root (`.`) DNS zone, has been targeted various times in the last decade [12, 16–18, 10], and even DNS providers have been victims of attacks [7], disrupting many of its domains [14]. DNS has been engineered with *layers of replication* to curb such attacks: first, a domain name may use multiple authoritative name servers. Second, each authoritative name server may employ IP anycast [8], which allows the same IP addresses to be replicated and announced from various locations, referred to as anycast sites. Third, each site, in turn, may locally use load balancers to distribute queries among multiple servers [10], increasing reliability even further.

Even though these measures are broadly employed, when domain names share the same DNS provider, they may be (unknowingly or not) sharing different levels of *infrastructure*, such as pipes, servers, and data centers. As many companies do not run their DNS infrastructure, instead of outsourcing to third-party DNS providers, identifying possible infrastructure sharing among many distinct domains, becomes a challenging endeavor. This infrastructure sharing may become a problem when a large enough DDoS attack takes place: if parts of the shared infrastructure become overwhelmed, all DNS zones under the service may experience problems too. As a consequence, many domains under zones may become unreachable. The Dyn attack [7] exemplifies the *collateral damage* when authoritative servers hosting multiple DNS zones are under attack.

The Internet DNS has been analyzed and studied by multiple authors [13, 11], yet few works focused on measuring the levels of the shared DNS infrastructure. Besides, measuring such levels of infrastructure sharing is a daunting task, given that researchers typically do not have access to DNS provider internals. As such, researchers have to resort to active measurements that allow to estimate, at the IP level, a certain degree of shared infrastructure, or analyze historical DNS datasets. This study has been done previously in some studies [3, 1, 4] that analyzed different aspects of the DNS ecosystem, such as the robustness and centralization of Top-Level Domains (TLD) [4] [3] and Root servers [1]. Despite shedding some light on infrastructure sharing at the TLD level by providing evidence that network-level infrastructure sharing is becoming more frequent over time. Those studies do not inspect DNS centralization on an Autonomous System (AS) level, derived from relying solely on third-party DNS providers. It is also possible that the same authoritative DNS server hosts multiple Fully Qualified Domain Names (FQDNs), and this shared infrastructed has not been analyzed.

Given this scenario, we introduce in this paper a methodology that allows measuring, to various degrees, the level of centralization of authoritative DNS servers using active DNS measurements. We focus our work on analyzing a possible centralization in authoritative DNS servers in the wild, for FQDNs. Also, we developed `dnstracker`, an opensource tool that implements our proposed methodology and provides a consolidated view of our findings. As a case study, we use `dnstracker` to analyze all domains of Alexa Top 1 Million [2] websites. We show that, in some cases, up to 12,000 authoritative DNS servers of the most visited websites share the same infrastructure of a DNS provider, and as such, could suffer from collateral damage in the event of an attack.

The remainder of this paper is organized as follows. In Section 2, we discuss the Related Work, reviewing previous efforts that analyzed DNS and its infrastructure. In Section 3, we describe the `dnstracker` methodology used to measure the DNS centralization and discuss its efficiency. In Section 4, we present our results. Finally, in Section 5, we conclude this work and discuss future directions.

2 Related Work

Moura *et al.* [1] analyzed the DDoS event suffered by the DNS Root servers in 2015. Between Nov. 30th to Dec. 1st, 2015, many of the Root DNS Letter Servers had an unusually high rate of a specific request, with a traffic rate a hundred times larger than the normal load. The authors highlighted that, even though these episodes did not target specific end-services, there was evidence of Internet services suffering from *collateral damage* because of sharing the DNS provider infrastructure with the DDoS target. In the 2015 attack, some `.nl` TLD servers were taken down as a side effect from the attack to the root DNS server. Even though that investigation provided a diagnosis of the events and highlighted some shreds of evidence of shared infrastructure, it did not investigate in depth the possible level of centralization in the DNS ecosystem.

Bates *et al.* [4] proposed a solution to measure how far the global DNS has preserved its distributed resilience, given the rise of cloud-based hosting and infrastructure. In their work, the authors analyzed the trends in concentration and diversification of the DNS infrastructure over time, where they sampled the 1,000 main US domains in the TLDs `.com`, `.net`, and `.org` according to Alexa Top Sites [2]. The authors also pointed out that their analysis focused on the traditional domains `.com`, `.net`, and `.org` because they are among the oldest TLDs, thus representing a broad spectrum of the current Internet. However, the authors recognize that their results might change if other TLDs, such as `.ru` and `.cn`, were taken into account. Despite providing some insight into the robustness of DNS, the work did not consider the possible concentration of authoritative DNS servers, which is a crucial point in infrastructure reliability. That, in turn, is covered by our work.

Allman *et al.* [3] carried out a study to observe the robustness of the DNS ecosystem. Their analysis was focused on Second-Level Domains (SLDs) (*e.g.*, , `icir.org`). In that study, the authors used two sets of zone files for the `.com`, `.net`, and `.org` TLDs. That data was collected over nine years. They performed an analysis of DNS infrastructure sharing. Initially, it was noted that 91% to 93% of the observed SLDs share, at least, one name server (by IP) with at worst one another SLD. In an approach based on individual SLDs, the authors observed that half of the SLDs share exactly one set of authoritative DNS servers with at the very least 163 other SLDs. Also, it was discovered that the largest group contains 9,000 SLDs that share the same set of authoritative DNS servers. In further analysis, by looking for shared infrastructure over IP blocks instead of single IPs, the authors found an even greater level of concentration. Besides, the authors point out that such network-level infrastructure sharing is becoming more common over time. Finally, they analyze the data to determine whether shared infrastructure occurs more frequently in domains with a higher or lower ranking. Their study, however, did not point to any general result or specific trends.

Considering the research carried out so far in the scientific community, there is strong evidence that suggests some level of DNS centralization. However, none of the works in the state-of-the-art has considered the centralization of authoritative DNS servers for FQDNs. Besides, it is also essential to have in mind not only the last hop but the hop before the last one, which is part of the contribution of our work. In Section 3, we describe our methodology to identify and quantify the centralization of the global DNS infrastructure.

3 `dnstracker`

The outsourcing trend of DNS providers poses several challenges to identifying possible infrastructure sharing. Also, the collateral damage cannot be directly assessed by a simple analysis of the IP addresses of different authoritative DNS servers. The different servers - each under their IP address block - could be hosted behind the same service provider's infrastructure. Moreover, due to the commercial nature of DNS providers, the data required to analyze such aggregation is rarely disclosed. An indication of this problem may be the presence of a common node (or single point of failure) in the routing path for a specific set of remote servers. For instance, if we get the IP address of `a.dns.nl`, one of the authoritative servers for the `.nl` Zone, and examine its AS, we will find it belongs to a DNS provider. NetNod runs this authoritative DNS server in Sweden. Hence, if other authoritative DNS servers are hosted on the same Netnod infrastructure, they start sharing the collateral damage in a potential DDoS attack. Below, we describe our proposed methodology for measuring DNS centralization, as well as the system we developed to implement it.

3.1 Measuring centralization

We propose an approach that identifies common points in routing paths. The approach presented is motivated by the fact that researchers do not have access to confidential company information. Therefore, analyzing single points of failure is the only way to infer possible collateral damage. This is essential information because the more domains that use a provider's infrastructure, the greater the visibility and, consequently, the higher the risk of attack and collateral damage.

Initially, for a given FQDN, we use a `dig` Linux command to uncover its authoritative DNS server. An FQDN has typically from two to four distinct authoritative DNS servers, but many domains share the same authoritative DNS server. Then, for every server uncovered with the dig command, we execute a custom `traceroute` command from a given vantage point. The command provides information about the addresses of each hop in the route from the vantage point to the domain's authoritative DNS server. Whenever a set of distinct servers, owned by different websites, are hosted behind the same provider's infrastructure, requests to these servers will share a common point - the network hop just before reaching its final destination, referred to as *Hop-Before-The-Last* (HBTL). If two different requests are served through a path whose HBTL is in the same AS, they are likely hosted by the same DNS provider, thus sharing the same infrastructure to some extent, as illustrated by Figure 1.

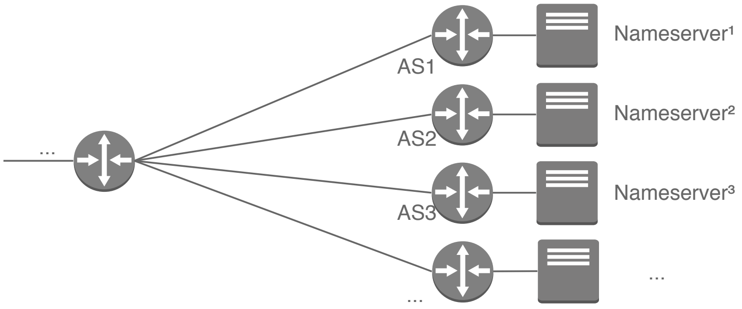

Fig. 1. Name Servers whose HBTL share the same AS infrastructure

From each ICMP Echo Reply obtained through *traceroute*, we extract the IPv4 address of each hop in the path to the authoritative DNS server. However, in our approach, we only store the relevant data of the last hop and the HBTL, as these are the most likely points of aggregation of infrastructure in the DNS ecosystem. For each IPv4 address, we use a publicly available BGP table [15] to obtain the corresponding AS of the hop, as well as the owner of the hop (*i.e.*, what company is responsible for running the infrastructure). We repeat this step for both the last hop and the HBTL. Listing 1.1, Listing 1.2 and Listing 1.3 present a step by step example of our methodology, using `bash` commands as examples of our proposed approach. For this example domain, we can see that authoritative DNS servers are hosted with the in-house infrastructure of UFRGS since both the ASes of the last hop and the HBTL are the same.

```
host $>dig ufrgs.br
;; AUTHORITY SECTION:
ufrgs.br 3600 IN NS ns1.ufrgs.br
ufrgs.br 3600 IN NS ns2.ufrgs.br
```

Listing 1.1. Uncovering domain authoritative DNS servers.

```
host $>traceroute ns1.ufrgs.br
...
19 ge-0-0-2-0.arn1-rndfw1.as1140.nl (94.198.152.11)
20 proteus.ufrgs.br (94.198.159.3)
```

Listing 1.2. Uncovering IPv4 addresses of last hop and HBTL

```
host $>ip2asn 94.198.152.11
AS Number: 1140, AS Description: UFRGS
host $>ip2asn 94.198.159.3
AS Number: 1140, AS Description: UFRGS
```

Listing 1.3. Mapping ASes of last hop and HBTL

Finally, having received the responses of all hops until the targeted authoritative DNS servers, and mapping the corresponding ASes of each hop, we store this information in our database for further analysis. When executed repeatedly, we can consolidate the millions of entries in the database to identify a possible aggregation of infrastructure in many different levels, as well as analyze the changes in the DNS ecosystem over time.

3.2 System Architecture

To support our methodology, the `dnstracker` tool was developed to collect DNS-related information and to expose the level of centralization of the DNS infrastructure. The source code for `dnstracker` is publicly available at GitHub[3]. Figure 2 presents the architecture of `dnstracker`.

`dnstracker` Agent. On the left side of the Figure 2, a group of `dnstracker` Agents retrieve information, using `traceroute`, from target authoritative DNS servers ①. The target servers are obtained from the list of the world's most popular websites, provided by Alexa Top 1 Million domain list open repository [2], accessed in January 2018 (hosted as a local conceptual database inside each agent). The agent applies our collection methodology for each domain in the list. After discovering information from all authoritative DNS servers, the `dnstracker` Agent exports the created datasets to the `dnstracker` Server ② using a REST API [6]. It is also important to mention that the traditional Linux implementation of `traceroute` does not support parallelism and, hence, it is not fit for actively crawling through large numbers of domains. We then implemented a custom version of `traceroute` with parallelism support, using Java, running inside each agent.

[3] https://github.com/ComputerNetworks-UFRGS/dnstracker

`dnstracker` Server. After exporting the collected data from the `dnstracker` agent, the `dnstracker` Server processes datasets to create appropriate visualization and exposes them for the system's users via HTTP ③. We used the Spring Boot v2.0.4 framework to prototype this Web user interface.

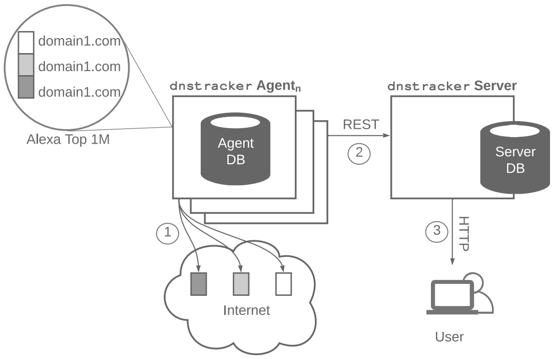

Fig. 2. `dnstracker` architecture

One of the benefits of `dnstracker` is that the tool automates several otherwise labor-intensive tasks. That includes tracing the route to all authoritative DNS servers in the database every month, identifying which ASes the last hop and HBTL belong to, as well consolidating the data to present possible aggregation. By using `dnstracker`, the user can observe various aspects of the DNS centralization. In the next Section, and through the use of `dnstracker`, we present the current picture of the Internet's DNS centralization.

4 Results

The application `dnstracker` that implements our methodology was deployed on two instances of Amazon EC2 in Sao Paulo, Brazil. One of these instances was used as a `dnstracker` Agent and the other as a `dnstracker` Server. These instances have the same configuration with a single-core CPU and RAM of 0.5GB. We executed our measurements several times a month, from January 2018 to May 2018, resulting in a dataset of millions of *traceroute* entries.

In this section, we present the results obtained through our proposed methodology. Having obtained the dataset through `dnstracker`, We focus on three facets to identify possible infrastructure centralization as well as identify the possible risk of collateral damage. First, we evaluate the concentration of authoritative servers per last-hop AS. Second, we measure the concentration of authoritative server ASes per HBTL AS. Third, we determine the total number of authoritative servers that shared the same HBTL. These three aspects enable us to measure the number of authoritative DNS servers that share AS infrastructure with other authoritative DNS servers, at both last hop level and

HBTL level. Finally, we analyze whether, among the top DNS providers, there is any growth trend in these aggregations over the measurement period. First, we describe our datasets. After that, we present our results.

4.1 Datasets

Table 1 presents a summary of the data collected throughout the five months of observations. Over the measurement period, the number of observed distinct authoritative DNS servers, ASes, and HBTL ASes remained stable. Additionally, the number of distinct authoritative DNS servers is much smaller than the number of investigated domains, since many domains share the same authoritative DNS servers. In fact, among the data we collected, one single server was authoritative for over 10,000 domains, belonging to a big DNS provider: DNSPod. This does not mean that the servers would lead to problems, since there may be multiple design measures in place to increase its the fault tolerance against DDoS attacks, but it also reveals the existence of actually shared infrastructure. In our samples, 136,421 out of the traced authoritative DNS servers had ASes in their routes that explicitly discarded ICMP echo requests, which hindered obtaining information about the HBTL of such server. Because of that, in the analysis of HBTL aggregation, such authoritative DNS servers were disregarded; circumventing this observation is a subject of future work in our research.

	DNS data			`traceroute` data	
Month	NS rec.	IPv4 (NS)	Last Hop ASes	HBTL IPv4	HBTL ASes
Jan - May	283,983	208,543	18,400	40,157	7,742

Table 1. Datasets generated by `dnstracker` for 2018 monthly measurements for Alexa 1 million domain names.

4.2 Aggregation of Authoritative DNS Server

First, we analyze the concentration of distinct authoritative DNS servers by last-hop ASes. As shown in Figure 3(a) "OVH, FR" is the provider with the most significant number of authoritative DNS servers in its infrastructure, aggregating more than 12,000 distinct authoritative DNS servers, each with multiple hosted domains. This means that, in case of a successful attack to this AS, over 77,419 websites would be unreachable, since clients would not be able to resolve FQDNs. Then, we can observe that the provider "HETZNER-AS, DE", holds 11,000 of the total authoritative servers, which represents 30,947 distinct websites hosted, followed by "UNIFIEDLAYER-AS-1 - Unified Layer, US", which hosts 6,000 authoritative servers behind his infrastructure representing 5,825 distinct websites. By analyzing the collected data, we can see that hosting an authoritative DNS server in these three providers presents a higher risk of suffering from collateral damage. These providers concentrate a large portion of the analyzed domains, and hence would likely be more target by attacks. In the other providers observed in Figure 3(a), we can see a margin ranging from 6,000 to 3,000 of the total authoritative DNS servers in each of the providers. These observations match with the results obtained in previous studies [3] by observing the IP blocks of the authoritative DNS servers. As indicated by previous work [4, 3], we also reinforce the presence of centralization in DNS services. We understand that large DNS providers such as these often offer multiple levels of redundancy to curb possible attacks. However, it is worth pointing out that DDoS attacks are becoming increasingly sophisticated, including o attack to the DynDNS infrastructure [7], so this should be a point of attention.

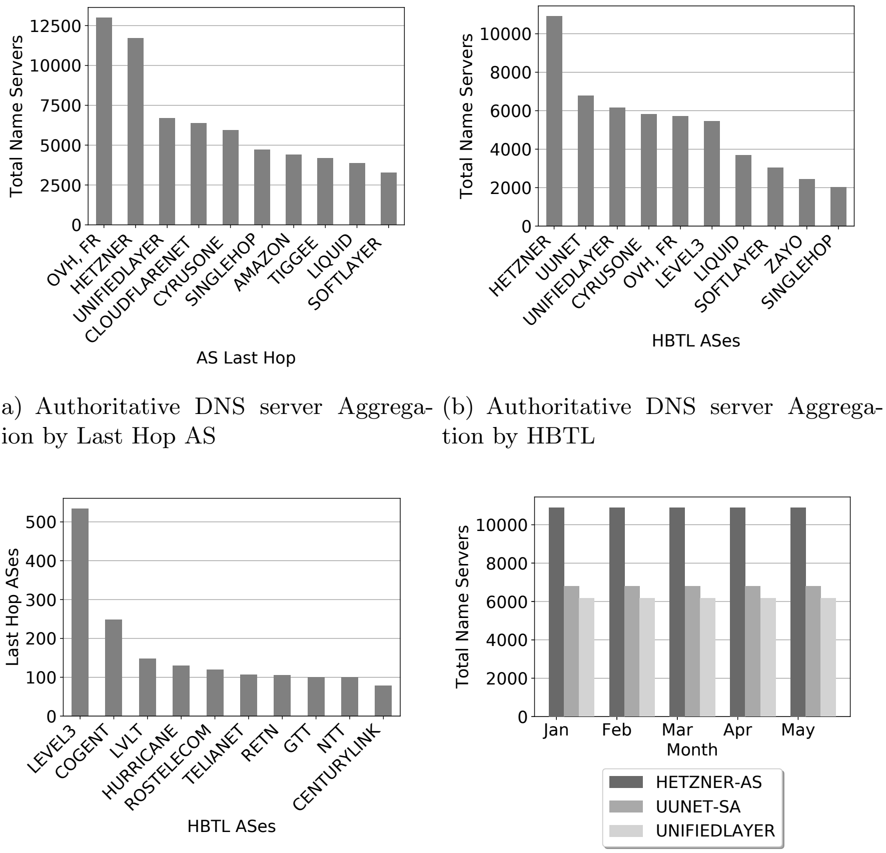

(a) Authoritative DNS server Aggregation by Last Hop AS

(b) Authoritative DNS server Aggregation by HBTL

(c) AS Aggregation by HBTL

(d) Authoritative DNS server aggregation by HTBL ASes over time

Fig. 3. Authoritative DNS server Aggregation

4.3 Aggregation of Authoritative DNS Servers by HBTL ASes

In addition to analyzing the aggregation of authoritative DNS servers in the last hop, we inspect the number of authoritative DNS servers that share the same HBTL, since it can be a single point of failure. Figure 3(b) present the top 10 HBTL ASes that aggregated more authoritative DNS servers.

The provider identified by "HETZNER-AS, DE" in Figure 3(b), shows that almost 11,000 of the total authoritative DNS servers share this same hop as its HBTL. We mention that HBTL may change depending on the vantage point. Other vantage points will be analyzed in future work. The "UUNET-SA - MCI Communications Services, Inc.", is shared by almost 7,000 authoritative servers as well. These numbers suggest the presence centralization in the DNS infrastructure itself, not only at the last hop, as mentioned by previous studies [4][3], but also in HBTL as well. Also, it is important to highlight once more that of these authoritative DNS servers resolve more than 100,000 domains. Hence, if an HBTL were to be taken down, hundreds of thousands of domains would become unreachable as collateral damage.

4.4 Aggregation of Last Hop ASes by HBTL ASes

Up to here, we focused on analyzing the concentration of authoritative DNS servers in each hop. However, when looking to third-party provider ASes, other services, such as storage, database, emails, maybe affected in addition to the hosted authoritative DNS servers. Hence, we also study the number of distinct ASes that share the same HBTL AS, aiming to identify points of shared network infrastructure that might represent the possibility of collateral damage to authoritative DNS servers, *i.e.*, an unrelated service might be targeted by an attack and still affect the DNS ecosystem because of shared infrastructure.

As we can see in Figure 3(c), in this assessment, the most noteworthy aggregation of last hop ASes occurs in the "LEVEL3 - Level 3 Communications, Inc., US" HTBL AS. Level 3 is one of the top infrastructure providers in the world, so this is a natural result. However, the number of last hop ASes that share its infrastructure is large, amounting to over 500 different ASes. The second-largest HBTL aggregation, provider "COGENT-174 - Cogent Communications, US", has less than half of Level 3's amount, with 200 AS behind it. Although the concentration of ASes behind a single hop has probably more to do with delivery structure than surely on DNS services, such a concentration increases the chance of problems for a more considerable amount of service if targeted by a large-scale attack.

4.5 Summary

The analysis we provided in the previous subsection showed a considerable amount of infrastructure sharing in many different levels of the DNS ecosystem. In particular, the level of concentration of authoritative DNS servers by

HBTL is worth highlighting, as most DNS operators are not aware of such concentration when contracting hosting services. Looking solely at the last hop ASes for concentration, on the other side, maybe misleading because many companies (*e.g.*, facebook.com) may advertise their ASes for their authoritative servers but still rely on the infrastructure of a third-party provider. In such cases, the possibility of collateral damage is still present, but undetected so far.

Lastly, as we evaluate DNS aggregation over five months, we must look at the difference in HBTL aggregation of authoritative servers during this period. By doing so, one may be able to identify if a trend of centralizing the DNS infrastructure exists, at least in such a period. Figure 3(d) presents the aggregation level of the top 3 HBTL ASes; over each month, we traced the authoritative servers. By observing the temporal graph, the centralization of authoritative servers of the Alexa Top 1 Million remained stable in the period. This is consistent with the general assumption that the DNS infrastructure is stable and robust. Also, that can be justified by the fact that the observed providers offer reliability, and there is no need for frequent changes in hosting them. However, this does not mean that there is no centralization trend considering a more significant time window.

5 Conclusions and Future Work

In this paper, we presented dnstracker, a tool to measure the centralization and shared infrastructure of Internet's DNS using active DNS measurements. dnstracker implements our proposed methodology that relies on traceroute to trace and get the informations about the last two hops of a authoritative DNS servers. We focus our work on analyzing centralization in the DNS infrastructure in the wild, for FQDNs. As a case study, we used dnstracker to analyze all domains of the Alexa Top 1 Million [2] websites. The analysis showed a considerable amount of infrastructure sharing, in many different levels of the DNS ecosystem. We show that, in some cases, up to 12,000 authoritative DNS servers of the most visited websites share the same infrastructure of big DNS providers, and thus could suffer from collateral damage in the event of an attack. In addition, we analyzed our measurements collected during 5 months to try to identify a centralization trend in the DNS global infrastructure. However, no trend was identified in the period of time we collected our data.

The future directions of our research include observing DNS centralization from different vantage points in the Internet; we want to understand how influential a vantage point is in our observation methodology. We also want to exploit the Ripe Atlas infrastructure to carry out our analysis. Finally, at a more theoretical perspective, we are working on a centralization metric that will help network operators find the more appropriate hosts for their DNS needs.

Acknowledgement

We thank CNPq for the financial support. This research has been supported by call Universal 01/2016 (CNPq), project NFV Mentor process 423275/2016-0.

References

1. Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In: Proceedings of the 2016 ACM on Internet Measurement Conference - IMC '16. No. November 2015 (2016). https://doi.org/10.1145/2987443.2987446, http://dl.acm.org/citation.cfm?doid=2987443.2987446
2. Alexa: Alexa Top 1 Million (Jan 2018), http://s3.amazonaws.com/alexa-static/top-1m.csv.zip
3. Allman, M.: Comments on DNS Robustness. In: ACM Internet Measurement Conference (Nov 2018), to appear
4. Bates, S., Bowers, J., Greenstein, S., Weinstock, J., Zittrain, J.: Evidence of decreasing internet entropy: The lack of redundancy in dns resolution by major websites and services. Tech. rep., National Bureau of Economic Research (2018)
5. Elz, R., Bush, R., Bradner, S., Patton, M.: Selection and Operation of Secondary DNS Servers. RFC 2182 (Best Current Practice) (Jul 1997). https://doi.org/10.17487/RFC2182, https://www.rfc-editor.org/rfc/rfc2182.txt
6. Fielding, R.T.: Architectural Styles and the Design of Network-based Software Architectures. Ph.D. thesis (2000), university of California, Irvine
7. Hilton, S.: Dyn analysis summary of Friday October 21 attack. Dyn blog https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/ (Oct 2016)
8. McPherson, D., Oran, D., Thaler, D., Osterweil, E.: Architectural Considerations of IP Anycast. RFC 7094 (Informational) (Jan 2014). https://doi.org/10.17487/RFC7094, https://www.rfc-editor.org/rfc/rfc7094.txt
9. Mockapetris, P.: Domain names - concepts and facilities. STD 13, Internet Engineering Task Force (November 1987)
10. Moura, G.C.M., de O. Schmidt, R., Heidemann, J., de Vries, W.B., Müller, M., Wei, L., Hesselman, C.: Anycast vs. DDoS: Evaluating the November 2015 root DNS event. In: Proceedings of the ACM Internet Measurement Conference (Nov 2016). https://doi.org/http://dx.doi.org/10.1145/2987443.2987446,
11. Mugali, A.A., Simpson, A.W., Walker, S.K., et al.: System and method for detecting dns traffic anomalies (Oct 27 2015), uS Patent 9,172,716
12. Paul Vixie and Gerry Sneeringer and Mark Schleifer: Events of 21-oct-2002 (Oct 2002), http://c.root-servers.org/october21.txt
13. Perdisci, R., Corona, I., Giacinto, G.: Early detection of malicious flux networks via large-scale passive dns traffic analysis. IEEE Transactions on Dependable and Secure Computing **9**(5), 714–726 (2012)
14. Perlroth, N.: Hackers used new weapons to disrupt major websites across U.S. New York Times p. A1 (Oct 22 2016), http://www.nytimes.com/2016/10/22/business/internet-problems-attack.html
15. RIPE Network Coordination Centre: RIPE Atlas, https://atlas.ripe.net/
16. Root Server Operators: Events of 2015-11-30 (Nov 2015), http://root-servers.org/news/events-of-20151130.txt
17. Root Server Operators: Events of 2016-06-25. Tech. rep., Root Server Operators (June 29 2016), http://www.root-servers.org/news/events-of-20160625.txt
18. Weinberg, M., Wessels, D.: Review and analysis of attack traffic against A-root and J-root on November 30 and December 1, 2015. In: DNS OARC 24 – Buenos Aires, Argentina. https://indico.dns-oarc.net/event/22/session/4/contribution/7 (April 2016)