DNS Negative Caching in the Wild

Lior Shafir
Tel Aviv University

Yehuda Afek
Tel Aviv University

Anat Bremler-Barr
The Interdisciplinary Center

Neta Peleg
The Interdisciplinary Center

Matan Sabag
The Interdisciplinary Center

1 INTRODUCTION

In this work we measure what percentage of DNS recursive resolvers perform negative caching in the wild. We deploy our own authoritative name server and harness thousands of RIPE Atlas [3] sites spread over the globe to perform repeated DNS queries for non-existing sub-domains of our authoritative domain.

The Domain Name System (DNS) system, the Internet phone book, is accessed each time a user establishes a connection, in order to translate the domain name to an IP address. The DNS infrastructure consists mostly of two types of servers, authoritatives and resolvers. The DNS directory itself is stored in the authoritative name servers, which are organized in a delegated hierarchical tree. In costrast, Recursive resolvers are responsible to perform the DNS resolution and act on behalf of clients that ask them to resolve their queries. DNS recursive resolvers rely heavily on their caches to improve efficiency and performances.

DNS negative caches store negative responses. The most common negative responses are called NXDOMAIN records (Non-Existent Domain), indicating that a particular domain does not exist. The duration of caching for a negative response is governed by the TTL value of the SOA (start of authority) record created by the authoritative server that responded with the negative response. While in the past negative caching was an optional part of the DNS specification, today it is not optional [1], but there are still networks that do not do DNS negative caching.

Negative caching has a significant impact on the DNS infrastructure for two main reasons:

Large proportion of DNS traffic. On one hand, negative caching is important to improve DNS performance. According to RFC 2308 [1], *A large proportion of DNS traffic on the Internet could be eliminated if all resolvers implemented negative caching. With this in mind negative caching should no longer be seen as an optional part of a DNS resolver.* Chen et al. [4] analyze real-life DNS traces and show that the NXDOMAIN traffic constitutes almost 40% of the traffic from the authoritative structure to the recursive resolvers, and only 6% of the traffic from the recursives to the stub resolvers in a network that does not use negative caching. This is probably because resolvers in their experiments did not respect the negative caching.

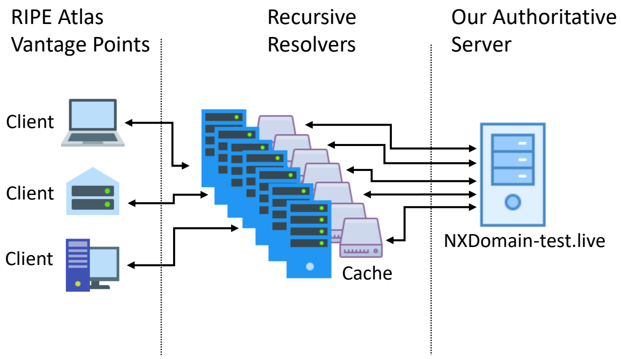

Figure 1: Measurement Architecture

NXDOMAIN attacks. On the other hand, negative caching might incur additional overhead and degrade performance during DDoS attacks. Recently, the DNS infrastructure has been an attractive target for different DDoS attacks, e.g., the notorious Mirai botnet [2]. As demonstrated by Mirai, such attack can be destructive to both authoritative and recursive DNS servers, disrupting many popular websites such as Twitter, Reddit, Netflix (and many others) and impacting millions of Internet users. Recent large scale attacks (NXDOMAIN attacks) were directly trying to take down parts of the DNS system by flooding the DNS servers with well-crafted queries that exploit unique vulnerabilities of the DNS infrastructure.

The NXDOMAIN attack (also called *water torture* [7], or Random Subdomain Attack, or Nonsense Name Attack) is a DDoS attack targeting DNS servers where the attacker issues many requests of randomly-generated non-existent sub-domains of the target domain (e.g., fake1.google.com, fake2.google.com). Since the sub-domains are random, their records are not present in the recursive resolvers caches, thus the malicious requests always reach the target authoritative server. Although some recursive implementations include a separate cache for NXDOMAIN responses (e.g., BIND), during such attacks the cache is filled with many NXDOMAIN records, until it overflows, which might lead to a critical performance degradation.

SIGCOMM '19, August 19–23, 2019, Beijing, China
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6886-5/19/08. . . $15.00
https://doi.org/10.1145/3342280.3342338

Here we measure and evaluate negative caching by performing several controlled experiments in which we issue a unique query to our authoritative server from each one of the thousands RIPE Atlas clients [3].

2 MEASUREMENT OVERVIEW

The experiment architecture is shown in Figure 1. We sent queries from 7,174 RIPE Atlas probes, spread all over the world. We placed an authoritative DNS server that responds to queries for our domains on an EC2 AWS machine in Ohio, USA. The authoritative machine is running Ubuntu OS, and BIND 9 (an open source DNS server) under authoritative mode. The authoritative zone file does not contain any A record, but only an SOA DNS record that specifies the negative TTL value (600 seconds).

We note that this measurement has two limitations: (i) The RIPE Atlas probes are spread unevenly over the world [6] (the majority of the probes (4632) are located in Europe), (ii) Some of the observed recursive resolvers are public resolvers. Moreover, we found that some resolvers use multiple caches [5], and appear as different IP addresses at the authoritative server. (Further analysis in §3).

Experiment Flow. With each probe we associate a unique non-existing sub-domain of our authoritative domain (e.g., probeID.nxdomain-test.live). From each probe we issue two queries 300 seconds apart. The first query should always reach the authoritative server since it is unique, thus not present in any of the recursive resolver caches. If both requests are observed by the authoritative server, the probe's DNS resolver is considered as a recursive server that does not perform negative caching. The experiment is repeated several times to verify the results consistency. Most of the probes were active in several experiments, however, a probe is considered as non-negative-caching only if negative caching was not observed in any experiment it was part of.

3 PRELIMINARY RESULTS

Our results show that 866 out of the 7,174 probes (12.07%) did not receive any cached response (all the requests reached the authoritative server, i.e., do not do negative caching).

Continent	No negative cache	Total #Probes	Percentage
Asia	134	900	14.89%
N. America	164	1144	14.34%
Oceania	21	178	11.80%
Africa	19	167	11.38%
Europe	513	4632	11.08%
S. America	15	152	9.87%

Table 1: Number of probes with no negative caching aggregated by continents.

The second column in Table 1 is the number of probes that did not receive any cached negative response in each continent. Asia has the highest percentage (14.89%) of measurements without NXDOMAIN caching, and S. America has the lowest (9.87%)

Inspecting the probes DNS requests that reached our authoritative servers reveals that a significant amount of probe requests arrive at the authoritative server from more than one source IP address as their DNS resolver, even in a single experiment. The reason is that public resolvers and many non-public DNS resolvers use load balancing techniques, with multiple caches. This phenomena skews our results, increasing the percentage of non-negative caching. Two different directions were taken to handle this skew in the measurement: first, to simply count only probes that appear to come from the same IP address (which is a resolver address) in both queries, however, this still leaves the case that more than one probe is behind the same resolver. Therefore, as a second direction we count resolvers rather than probes. In the resolvers we should consider the case in which a resolver has multiple caches, and/or multiple ingress and/or egress IP addresses.

Single IP probes. Notice that if a single probe was active in multiple experiments, we considered only those in which it arrives from a single IP address. Excluding any probe that arrives to the authoritative through different IP addresses leaves us with 4,774 probes (66%). Only 185 of them (3.88%) have no negative caching. Table 2 summarizes this single-IP probe experiment.

Continent	No negative cache	#Probes (single IP addr.)	Percentage
S. America	7	112	6.25%
Europe	130	3080	4.22%
Africa	4	119	3.36%
N. America	24	740	3.24%
Asia	17	596	2.86%
Oceania	3	127	2.36%

Table 2: Number of probes (observed with a single IP address) with no negative caching aggregated by continents.

Counting Resolvers. We aggregated all the IP addresses observed in the authoritative server into /24 subnets assuming each is a resolver. To determine whether a resolver does not perform negative caching, we verified that none of its associated probes experienced negative caching (less than two requests reached the authoritative server). This aggregation resulted in 3309 different resolvers, from which 259 were observed with no negative caching (7.83%).

4 CONCLUSIONS AND FUTURE WORK

Our results show that some resolvers still operate with no negative caching. The gap between the results over all the probes (12.07%, see Table 1) and the results over probes whose resolvers were not observed with multiple IP addresses (3.88%, see Table 2) indicates that many resolvers (both ISP and public resolvers) use load balancing techniques. Thus consecutive requests are likely to come from different servers even if the requests came from within the same subnet. As part of our ongoing and future work, we focus on analyzing the impact of public resolvers and ISP resolvers that use multiple caches on our results and on DNS caching.

REFERENCES

[1] Andrews, M. Negative caching of dns queries (dns ncache). RFC 2308, RFC Editor, March 1998.

[2] Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., Cochran, J., Durumeric, Z., Halderman, J. A., Invernizzi, L., Kallitsis, M., et al. Understanding the mirai botnet. In *26th USENIX Security Symposium USENIX Security 17)* (2017), pp. 1093–1110.

[3] Atlas, R. Ripe network coordination centre, May 2019.

[4] Chen, Y., Antonakakis, M., Perdisci, R., Nadji, Y., Dagon, D., and Lee, W. Dns noise: Measuring the pervasiveness of disposable domains in modern dns traffic. In *DSN* (2014), IEEE Computer Society, pp. 598–609.

[5] Klein, A., Shulman, H., and Waidner, M. Counting in the dark: Dns caches discovery and enumeration in the internet. In *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)* (2017), IEEE, pp. 367–378.

[6] Moura, G. C. M., Heidemann, J., Müller, M., de O. Schmidt, R., and Davids, M. When the dike breaks: Dissecting dns defenses during ddos. In *Proceedings of the Internet Measurement Conference 2018* (New York, NY, USA, 2018), IMC '18, ACM, pp. 8–21.

[7] Secure64:. Water torture, a slow drip dns ddos attack, Feb. 2014.