

SIDN Labs

<https://sidnlabs.nl>

May 31st, 2018

Technical Report

Title: When the Dike Breaks: Dissecting DNS Defenses During DDoS

Authors: Giovane C. M. Moura, John Heidemann, Moritz Müller, Ricardo de O. Schmidt, and Marco Davids

Technical Report: ISI-TR-725

Citation:

- Giovane C. M. Moura, John Heidemann, Moritz Müller, Ricardo de O. Schmidt and Marco Davids. **When the Dike Breaks: Dissecting DNS Defenses During DDoS (extended)**. Technical Report ISI-TR-725. USC/Information Sciences Institute.

- Bibtex:

```
@techreport{Moura18a,
  author = {Moura, Giovane C. M. and Heidemann, John and M{"u"}ller,
    Moritz and de O. Schmidt, Ricardo and Davids, Marco},
  title = {When the Dike Breaks: Dissecting {DNS}
    Defenses During {DDoS} (extended)},
  institution = {USC/Information Sciences Institute},
  year = {2018},
  sortdate = {2018-05-30},
  number = {ISI-TR-725},
  month = may,
  keywords = {anycast, dns, ddos, root ddos},
  url = {https://www.isi.edu/%7ejohnh/PAPERS/Moura18a.html},
  pdfurl = {https://www.isi.edu/%7ejohnh/PAPERS/Moura18a.pdf},
  otherurl = {ftp://ftp.isi.edu/isi-pubs/tr-709.pdf},
  copyrightholder = {authors}
}
```

When the Dike Breaks: Dissecting DNS Defenses During DDoS (extended)

USC/ISI Technical Report ISI-TR-725

May 2018

Giovane C. M. Moura
SIDN Labs and TU Delft

John Heidemann
USC/Information Sciences
Institute

Moritz Müller
SIDN Labs and University of
Twente

Ricardo de O. Schmidt
University of Passo Fundo

Marco Davids
SIDN Labs

ABSTRACT

The Internet’s Domain Name System (DNS) is a frequent target of Distributed Denial-of-Service (DDoS) attacks, but such attacks have had very different outcomes—some attacks have disabled major public websites, while the external effects of other attacks have been minimal. While on one hand the DNS protocol is a relatively simple, the *system* has many moving parts, with multiple levels of caching and retries and replicated servers. This paper uses controlled experiments to examine how these mechanisms affect DNS resilience and latency, exploring both the client side’s DNS *user experience*, and server-side traffic. We find that, for about 30% of clients, caching is not effective. However, when caches are full they allow about half of clients to ride out server outages. Caching and retries together allow up to half of the clients to tolerate DDoS attacks that result in 90% query loss, and almost all clients to tolerate attacks resulting in 50% packet loss. While clients may get service during an attack, tail-latency increases for clients. For servers, retries during DDoS attacks increase normal traffic up to 8×. Our findings about caching and retries help explain why users see service outages from real-world DDoS events, but minimal visible effects from others.

KEYWORDS

DNS, recursive DNS servers, caching

1 INTRODUCTION

DDoS attacks have been growing in frequency and intensity for more than a decade. Large attacks have grown from 100 Gb/s in 2012 [4] to over 1 Tb/s in 2017 [30], and 1.7 Tb/s in 2018 [16, 21]. Such attacks are sourced from large botnets (for example, with Mirai peaking at 600k hosts [3]), fueled by the continued deployment of new devices. Gigabit-size attacks are commodities today, selling for a few dollars via DDoS-as-a-Service [40].

The Internet’s Domain Name System (DNS) is a popular target of DDoS attacks. DNS is a very visible target, since name resolution is a necessarily step in almost any Internet activity. Root DNS servers have seen multiple attacks over more than a decade [22, 29, 37, 38, 48], as well as threats of attacks [44]. Other authoritative DNS servers have also been attacked, with the huge October 2016 against Dyn [14] resulting in disruptions at a number of prominent websites, including Twitter, Netflix and the New York Times [30].

The *outcome* of these attacks on services has varied considerably. The October 2016 Dyn attack is noted for disruption to websites that were using Dyn as their DNS provider, and extortion attempts often include DDoS [31]. However, multiple attacks on the DNS Root have occurred with, as far as has been reported, no visible service outages [37, 38].

An important factor in DNS resilience is heavy use of caching—we believe that differences in use of DNS caching contribute to the very different outcomes when DNS is subject to DDoS attack. Yet understanding DNS caching is difficult, with requests traveling from *stub resolvers* in web browsers and at client computers, to *recursive resolvers* at ISPs, which in turn talk to multiple *authoritative DNS servers*. There are many parts involved to fully resolve a DNS name like www.example.com: while the goal is an IP address (an A or AAAA DNS record), multiple levels of the hierarchy (root, .com, and .example.com) are often on on different servers (requiring NS records), and DNSSEC may require additional information (RRSIG, DNSKEY, and DS records). Each of these records may have different cache lifetimes (TTLs), by choice of the operator or because of DNS cache timeouts. We explore caching through controlled experiments (§3) and analysis of real-world use (§4).

Another factor in DNS resilience is recursives that retry queries when do not receive an answer. Recursives fail to receive answers occasionally due to packet loss, but pervasively during a DDoS attack. We examine how retries interact

with caching to mitigate DDoS attacks for loss during DDoS attacks (§5) and their effects on authoritatives (§6).

This paper assesses DNS resilience during DDoS attacks, with the goal of explaining different outcomes from different attacks (§8) through understanding the role of DNS caching, retries, and use of multiple DNS recursive resolvers. It is common knowledge that these factors “help”, but knowing *how* and *how much* each contributes builds confidence in defenses. We consider this question both as an operator of an authoritative server, and as a user, defining the *DNS user experience* latency and reliability users should expect.

Our first contribution is to build an end-to-end understanding of DNS caching. Our key result is that *caching often behaves as expected, but about 30% of the time clients do not benefit from caching*. While prior work has shown DNS resolution infrastructure can be quite complex [43], we establish a baseline DNS user experience by assessing the prevalence of DNS caching in the “wild” through both active measurements (§3) and through analysis of passive data from two DNS zones (.nl and the root zone §4).

Our second contribution is to show that *DNS mechanisms of caching and retries provide significant resilience client user experience during denial-of-service (DDoS) attacks* (§5). For example, about half of the clients continue to receive service during a full outage when caches are primed. Often DDoS attacks cause very high loss but not a complete outage. Even with very heavy query loss (90%) on all authoritatives, full caches protect half of the clients, and retries protect 30%. With a DDoS that causes 50% packet loss, nearly all clients succeed, although with greater latency than typical.

Third, we show that there is a large increase in legitimate traffic during DDoS attacks—up to 8× the number of queries (§6). While DNS servers are typically heavily over-provisioned, this result suggests the need to review by how much. It also shows the importance that stub and recursive resolvers follow best practices and exponentially back-off queries after failure so as to not add fuel to the DDoS fire.

Our final contribution is to suggest why users have seen relatively little impact from root servers DDoSes, while customers from some DNS providers quickly felt attacks (§8). When cache lifetimes are longer than the duration of a DDoS attack, many clients will see service for names popular enough to be cached. While many websites use short cache timeouts to support control with DNS-based load balancing, they may wish to consider longer timeouts as part of strategies for DDoS defense. Retries provide additional coverage, preventing failures during large attacks.

All public datasets from this paper is available [23], with our RIPE Atlas data also available from RIPE [34]. Privacy concerns prevent release of .nl and Root data (§4).

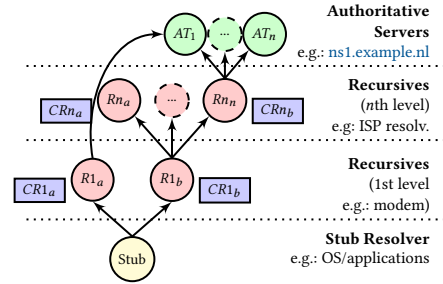


Figure 1: Relationship between stub resolver (yellow), recursive resolvers (red) with their caches (blue), and authoritative servers (green).

2 BACKGROUND

As background, we briefly review the components of the DNS ecosystem and how they interact with IP anycast.

2.1 DNS Resolvers: Stubs, Recursives, and Authoritatives

Figure 1 shows the relationship between three components of DNS resolvers: stubs and recursives resolvers and authoritative servers. Authoritative servers (authoritatives hereafter) are servers that know the contents of a given DNS zone and can answer queries without asking other servers [11].

Resolvers on the other hand, are servers that can ask, on behalf of others, queries to other servers [19]. *Stub* resolvers run directly on clients and query one or a few *recursive* resolvers (shortened to stubs and recursives here). Recursives perform the full resolution of a domain name, querying one or more authoritatives, while caching responses to avoid repeatedly requesting popular domains (e.g., .com or .google.com). Sometimes recursives operate in multiple tiers, with clients talking directly to R1 resolvers, that forward queries to other Rn resolvers, that ultimately contact authoritatives.

In practice, stubs are part of the client OS or browser, recursives are provided by ISPs, and authoritatives are run by DNS providers or large organizations. Multi-level recursives might have R1 at a home router and Rn in the ISP, or might occur in large, public DNS providers.

2.2 Authoritative Replication and IP Anycast

Replication of a DNS service is important to support high reliability and capacity and to reduce latency. DNS as two complementary mechanisms to replicate service. First, the protocol itself supports *nameserver replication* of DNS service for a zone (.nl or example.nl), where multiple servers operate on different IP addresses, listed by that zone’s NS records. Second, each of these servers can run from multiple physical locations with *IP anycast* by announcing the same IP address from each and allowing Internet routing




Figure 2: Stub resolver and Anycast Recursive

(BGP) to associate clients with each anycast site. Nameserver replication is recommended for all zones, and IP anycast is used by most large zones such as the DNS Root and most top-level domains [22, 39]. IP anycast is also widely used by *public resolvers*, recursive resolvers that are open for use by anyone on the Internet, such as Google Public DNS [12], OpenDNS [25], Quad9 [33], and 1.1.1.1 [1].

Figure 2 illustrates the relationship of elements in the DNS infrastructure when anycast is in place. The recursive resolver R is announced using an anycast prefix from five different *anycast sites* ($R1$ to $R5$). Each anycast site can have multiple servers, with DNS traffic designated to them by load balancing algorithms.

When the stub resolver sends a DNS query to the anycast address of R , BGP forwards the query to the nearest site ($R3$ in this example) of R —note that “nearest” metrics can vary [9]. The BGP mapping between source and anycast destination is known as anycast catchment, which under normal conditions is very stable across the Internet [47]. On receiving the query, $R3$ contacts the pertinent authoritative AT to resolve the queried domain name. The communication from $R3$ to AT is done using $R3$ ’s unicast address, ensuring the reply from AT is sent back to $R3$ instead of any other anycast site from R . On receiving the answer from AT , $R3$ replies to the stub resolver with the answer to its query; and for this reply $R3$ uses its anycast address as source.

2.3 DNS Caching with Time-to-Live (TTLs)

DNS depends on caching to reduce latency to users and load on servers. Authoritatives provide responses that are then cached in applications, stub resolvers, and recursive resolvers. We next describe its loose consistency model.

An authoritative resolver defines the lifetime of each result by its *Time-to-Live* (TTL); although TTLs is not usually exposed to users, this information is propagated through recursive resolvers.

Once cached, results cannot be invalidated directly; the only way to refresh a response with a new result is to wait for the TTL to expire. Operators therefore set TTLs careful. Content delivery networks (CDNs) often use DNS to steer users to different content servers. They therefore use very short TTLs (60 seconds or less) to force clients to re-query

frequently, providing opportunities to redirect clients with DNS in response to changes in load or server availability [26]. Alternatively, DNS data for top-level domains often has TTLs of hours or days. Such long TTLs reduce latency for clients (the reply can be reused immediately if it is in the cache of a recursive resolver) and reduce load on servers for commonly used top-level domains and slowly changing DNSSEC information.

3 DNS CACHING IN CONTROLLED EXPERIMENTS

To understand the role of caching at recursive resolvers in protection during failure of authoritative servers, we first must understand *how often are cache lifetimes (TTLs) honored*.

In the best-case scenario, authoritative DNS operators may expect clients to be able to reach domains under their zones even if their authoritative servers are unreachable, for as long as cached values in the recursive resolvers remain “valid” (*i.e.*, TTL not expired). Given the large variety of recursive implementations, we pose the following question: *from a user point-of-view, can we rely on recursive caching when authoritatives fail?*

To understand cache lifetimes in practice, we carry out controlled measurements from thousands of servers. These measurements determine how well caches work in the field, complementing our understanding of how open source implementations work from their source code. This study is important because operational software can vary and large deployments often use heavily customization or closed source implementations [43].

3.1 Potential Impediments to Caching

Although DNS records should logically be cached for the full TTL, a number of factors can shorten cache lifetimes in practice: caches are of limited size, caches may be flushed prematurely, and large resolvers may have fragmented caches. We briefly describe these factors here; understanding how often they occur motivates the measurements we carry out.

Caches are of limited size. Unbound, for example, defaults to a 4 MB limit, but the values are configurable. In practice, DNS results are small enough and caches large enough that cache sizes are usually not a limiting factor. Recursive resolvers may also override record TTLs, imposing either a minimum or maximum value [46].

Caches can be flushed explicitly (at the request of the cache operator), or accidentally on restart of the software or reboot of the machine running the cache.

Finally, some recursive resolvers handle very high request rates—consider a major ISP or public resolver [12, 25, 33]. Large recursive resolvers are often implemented as many separate recursive resolvers behind a load balancer or on IP anycast. In such cases the caches are fragmented as each machine has

TTL	60	1800	3600	86400	3600-10min
Probes	9173	9216	8971	9150	9189
Probes (val.)	8725	8788	8549	8750	8772
Probes (disc.)	448	428	422	400	417
VPs	15330	15447	15052	15345	15397
Queries	94856	96095	93723	95780	191931
Answers	90525	91795	89470	91495	183388
Answer (val.)	90079	91461	89150	91172	182731
Answers (disc.)	446	334	323	323	657

Table 1: Caching baseline experiments [34].

its own, independent cache. In practice these will reduce the cache hit rate.

3.2 Measurement Design

To evaluate caching we use controlled experiments where we query from specific names to authoritative servers we run from thousands of RIPE Atlas sites. Our goal is to measure whether the TTL we define for the RRs of our controlled domain is honored across recursives.

Authoritative servers: we deploy two authoritatives that answer for our new domain name (cachetest.nl). We place the authoritatives on virtual machines in the same datacenter (Amazon EC2 in Frankfurt, Germany), each at a distinct unicast, IPv4 addresses. Each authoritative runs BIND 9.10.3. Since both authoritatives are in the same datacenter, they will have similar latencies to recursives, so we expect recursives to evenly distribute queries between both authoritative servers [24].

Vantage Points: We issue queries to our controlled domain from around 9k RIPE Atlas probes [35]. Atlas Probes are distributed across 3.3k ASes, with about one third hosting multiple VPs. Atlas software causes each probe to issue queries to each of its local recursive resolvers, so our *vantage points* (VPs) are the tuple of probe and recursive. The result is that we have more than 15k VPs (Table 1).

Queries and Caching: We take several steps to ensure that caching does not interfere with queries. First, each *query* is for a name unique to the probe: each probe requests an AAAA record for `{probeid}.cachetest.nl`, where `{probeid}` is the probe’s the unique identifier. Each *reply* is also customized. In the AAAA reply we encode three fields that are used to determine the effectiveness of caching (§3.4). Each IPv6 address in the answer is the concatenation of four values (in hex):

- prefix** is a fixed, 64-bit value (fd0f:3897:faf7:a375)
- serial** is a 8-bit value, incremented every 10 minutes (zone file rotation), allowing us to associate replies with specific query rounds
- probeid** is the unique Atlas probeID [36] encoded in 8 bits, to associate the query with the reply

tll is a 16-bit value of the TTL value we configure per experiment

We increment the serial number in each AAAA record and reload the zone (with a new zone serial number), every 10 minutes. The serial number in each reply allows us to distinguish cached results from prior rounds from fresh data in this round.

Atlas DNS queries timeout after 5 seconds, reporting “no answer”. We will see this occur in our emulated DDoS events.

For example, a VP with probeID 1414 shall send a DNS query for AAAA record for the domain name 1414.cachetest.nl, and should receive AAAA answer in the format `$PREFIX:1:586::3c`, where `$SERIAL=1` and `$TTL=60`.

We focus on DNS over UDP on IPv4, not TCP or IPv6. We use send only IPv4 queries from Atlas Probes, and serve only IPv4 authoritatives, but the IPv6 may be used inside multi-level recursives. Our work could extend to cover other protocols, but we did not want to complicate analysis the the orthogonal issue of protocol selection. . We focus on DNS over UDP because it is by far the dominant transport protocol today (more than 97% of connections for `.nl` [45] and most Root DNS servers [15]).

Atlas’ Geographic distribution: It is well known that the global distribution of RIPE Atlas probes is uneven; Europe has far more than elsewhere [5, 6, 41]. Although quantitative data analysis might be generally affected by this distribution bias, our qualitative analysis, contributions and conclusions do not depend on the geographical location of probes.

3.3 Datasets

We carried out five experiments, varying the cache lifetime (TTL) and probing frequency from the VPs. Table 1 lists the parameters of experiments. In the first four measurements, the probing interval was fixed to 20 minutes, and TTL for each AAAA was set to 60, 1800, 3600 and 86400 seconds, all frequently used TTL values. For the fifth measurement we fixed the TTL value to 3600 seconds, and reduced the probing interval to 10 minutes to get better resolution of dynamics.

In each experiment, queries were sent from about 9k Atlas probes. Out of these, 400-448 were disregarded given they either do not return an answer (Probes disc.). Ultimately, each experiment was carried with 15k VPs when considering each Atlas probe’s recursive resolvers. A few queries do not receive valid DNS answers; for example, answers with various DNS error codes (for example, SERVFAIL and REFUSED [20]). We discard these replies from our dataset and consider only successful replies. (Around 3.5% to 4.9% of answers are discarded across measurements due to errors or the return of NS records instead of the desired AAAA.)

TTL	60	1800	3600	86400	3600-10m
Answers (valid)	90079	91461	89150	91172	182731
1-answer VPs	38	51	49	35	17
Warm-up (AAi)	15292	15396	15003	15310	15380
Duplicates	25	23	25	22	23
Unique	15267	15373	14978	15288	15357
TTL as zone	14991	15046	14703	10618	15092
TTL altered	276	327	275	4670	265
AA	74435	21574	10230	681	11797
CC	235	29616	39472	51667	107760
CCdec.	4	5	1973	4045	9589
AC	37	24645	24091	23202	47262
TTL as zone	2	24584	23649	13487	43814
TTL altered	35	61	442	9715	3448
CA	42	179	305	277	515
CAdec.	7	3	21	29	65

Table 2: Valid DNS answers (expected/observed)

Overall, the ~9k probes sent, 20 minute pacing ($T = 20$) results in about 93-96k queries to cachetest.nl, and measurements with 10 minute pacing result in about double that many queries.

3.4 TTL distribution: expected vs. observed

We next investigate how often recursive resolvers honor the full TTL provided by authoritative servers. Our goal is to classify the valid DNS answers from Table 1 into four categories, based on where the answer comes from, and where we *expect* it to come from:

- AA answers expected and correctly from the authoritative
- CC expected and correct from a recursive cache (cache hits)
- AC answers from the authoritative, but expected to be from the recursive’s cache (a cache miss)
- CA answers from a recursive’s cache, but expected from the authoritative (an extended cache)

To determine if a query should be answered by the cache of the recursive, we track the state of prior queries and responses, and the estimated TTL. Tracking state is not hard since we know the initial TTL and all queries to the zone, and we encode the serial number and the TTL in the AAAA reply (§3.2).

Cold Caches and Rewriting TTLs: We first consider queries made against a cold cache (the first query of a unique name) to test how many recursives *override* the TTL. We know that this happens at some sites, such as Amazon EC2 where the default recursive resolver caps all TTLs to 60 s [32].

Table 2 shows the results of our five experiments, in which we classify the valid answers from Table 1. Before classifying them, we first disregard VPs that had only one answer (1-answer VPs) since we cannot evaluate their caches status with one answer only (maximum 51 VPs out of 15,000 for




Figure 3: Classification of subsequent answers with warm cache

the experiments). Then, we classify the remaining queries as Warm-up queries AA_i , all of which are type AA (expected and answered by the authoritative server).

We see some duplicate responses; for these we use the timestamp of the very first AA_i received. We then classify each unique AA_i by comparing the TTL value returned by the recursive with the expected TTL that is encoded in the AAAA answer (fixed per experiment). The *TTL as zone* line counts the answers we expect to get, while *TTL altered* shows that a few hundred recursive resolvers alter the TTL. If these two values differ by more than 10%, we report TTL altered.

We see that the vast majority of recursives honor small TTLs, with about 2% truncating them (275 to 327 of about 15000, depending on the experiment’s TTL). We and others (§7) see truncation from multiple ASes. The exception is for queries with day-long TTLs (86400 s), where 30% of recursives shorten it. We conclude that wholesale TTL shortening does not occur, at least for TTLs of an hour or less.

TTLs with Warm Cache: We next consider a warm cache—subsequent queries where we believe the recursive should have the prior answer cached and classify them according to the proposed categories (AA, CC, AC, and CC).

Figure 3 shows a histogram of this classifications (numbers shown on Table 2). We see that most answers we receive show expected caching behavior. For 60 s TTLs (the left bar), we expect no queries to be cached when we requery 20 minutes (1200 s) later, and we see no cache hits. We see only a handful of CA-type replies, where we expect the authoritative to reply and the recursive does instead. We conclude that under normal operations (with authoritatives respond), recursive resolvers do not serve stale results (as has been proposed when the authoritative cannot be reached [17]).

For longer TTLs we see cache misses (AC responses) fractions of 28 to 33%. information from recursives caches. Most of the AC answers did *not* alter the TTL (AC-over), *i.e.*, the cache miss was not due to TTL manipulations (Table 2). We do see many TTL modifications (about 42%) when the TTL

TTL	60	1800	3600	86400	3600-10m
AC Answers	37	24645	24091	23202	47,262
Public R_1	0	12000	11359	10869	21955
Google Public R_1	0	9693	9026	8585	17325
other Public R_1	0	2307	2333	2284	4630
Non-Public R_1	37	12645	12732	12333	25307
Google Public R_n	0	1196	1091	248	1708
other R_n	37	11449	11641	12085	23599

Table 3: AC answers public resolver classification.

is 1 day TTLs (86400 s). Some recursives may have an upper limit for TTL, and one day is considered a long TTL.

We conclude that DNS caches are fairly effective, with cache hits about 70% of the time. This estimate is likely a lower bound: we are the only users of our domain, and popular domains would see cache hits due to requests from other users. We only see TTL truncation for day-long TTLs. This result will help us understand the role of caching when authoritatives are under stress.

3.5 Public Recursives and Cache Fragmentation

Although we showed that most requests are cached as expected about 30% are not. We know that many DNS requests are served by public recursive resolvers today, several of which exist [1, 12, 25, 33]. We also know that public recursives often use anycast and load balancing [43] and that that can result in caches that are fragmented (not shared) across many servers. We next examine how many cache misses (type AC replies) are due to public recursives.

Although we control queriers and authoritative servers, there may be multiple levels of recursive resolvers in between. From Figure 1, we see the querier’s first-hop recursive (R_1) and the recursive that queries the authoritative (R_n). Fortunately, queries and replies are unique, so we can relate queries to the final recursive knowing the time (the query round) and the query source. For each query q , we extract the IP address of R_n and compare against a list of IP addresses for 96 public recursives (Appendix B) we obtain from DuckDuckGo search for “public dns” done on 2018-01-15.

Table 3 reexamines the AC replies from Table 2. With the exception of the measurements with TTL of 60 s, nearly half of AC answers (cache misses) are from queries to public R_1 recursives, and about three-quarters of these are from Google’s Public DNS. The other half of cache misses start at non-public recursives, but 10% of these eventually emerge from Google’s DNS.

Besides identifying public recursives, we also see evidence of cache fragmentation in answers from caches (CC and CA). Sometimes we see serial numbers in consecutive answers decrease. For example, one VP reports serial numbers 1, 3, 3,

7, 3, 3, suggesting that it is querying different recursives, one with serial 3 and another with serial 7 in its cache. We show these occurrences in Table 2 as CCdec. and CAdec. With longer TTLs we see more cache fragmentation, with 4.5% of answers showing fragmentation with day-long TTLs.

From these observations we conclude that cache misses result from several causes: (1) use of load balancers or anycast where servers lack shared caches, (2) first-level recursives that do not cache and have multiple second-level recursives, and (3) caches that clear between or somewhat long probing interval (10 or 20 minutes). Causes (1) and (2) occur in public resolvers (confirmed by Google [12]) and account for about half of the cache misses in our measurements.

4 CACHING IN A PRODUCTION ZONE

In §3 we showed that about one-third of queries to not confirm with caching expectations, based on controlled experiments to our test domain. We next examine this question for .nl, the country code domain (ccTLD) for the Netherlands and the Root (.) DNS zone. With traffic from “the wild” and a measurement target used by millions, this section uses a domain popular enough to stay in-cache at recursives.

4.1 Requests at .nl’s Authoritatives

We next apply this methodology to data for the .nl country-code top-level domain (ccTLD).

Methodology: We use passive observations of traffic to the .nl authoritative servers. We extract the source IP address from A-record queries for the domains ns[1-5].dns.nl. We analyze caching of referrals in Appendix C.

For each target name in the zone and source (some recursive server, identified by IP address), we build a timeseries of all requests and compute their interarrival time, Δ . Following the classification from §3.4, we label queries as: **AC** if $\Delta < TTL$, showing an unnecessary query to the authoritative; **AA** if $\Delta = TTL$, an expected cache refresh; and **CA** if $\Delta > TTL$, a delayed cache refresh. (We do not see cache hits and so there are no **CC** events.)

Dataset: At the time of our analysis (February 2018) there were 8 authoritative servers for the .nl zone. We collect traffic for the 4 unicast and one anycast authoritative servers, and store the data in ENTRADA [49] for analysis.

Since our data for .nl is incomplete, and we know recursives will query all authoritatives over time [24], our analysis represents a conservative estimate of TTL violations—we expect to miss some CA-type queries from resolvers to non-monitored authoritatives.

We collect data for a period of six hours on 2018-02-22 starting at 12:00 UTC. We only evaluate recursives that sent at least five queries for our domains of interest, omitting infrequent recursives (they do not change results noticeably). We discard duplicate queries, for example, a few retransmissions




Figure 4: ECDF of Δt for recursives with at least 5 queries to `ns1-ns5.dns.nl` (TTL of 3600 s.)

(less than 0.01% of the total queries). In total, we consider more than 485k queries from 7,779 different recursives.

Results: Figure 4 shows the Δt distributions we observed in our measurements. We see many queries near the origin, and peaks at 3600 s, 300 s, and 600 s.

About 28% of queries are very frequent, with an inter-arrival less than 10 s, and 32% of these are sent to multiple authoritatives. We believe these are due to recursives submitting queries in parallel to speed replies. Since these closely-timed queries are not related to recursive caching, we exclude them from analysis. The remaining data is 348k queries from 7,703 different recursives.

The largest peak is at 3600 s, what was expected: the name was queried and cached for the full hour TTL, then the next request causes the name to be re-fetched. These queries are all of type AA.

The several smaller peaks around 300 s and 600 s, as well as queries with other times less than 3600 s, correspond to type AC-queries—queries that could have been supplied from the cache but were not. These represent about 80% of the queries, with at least 46% of recursives sending one of these.

These AC queries occur because of TTL limiting, cache fragmentation, or other reasons that clear the cache. When we examine specific recursives, we see that about 30% are from a few resolvers of a Russian ISP. Only about 0.25% of queries are from Google’s public DNS, because `.nl` is often already cached.

We conclude that measurements of popular domains within `.nl` show that about 30% of queries are from recursives that honor the full TTL, but most queries will be more frequent.

4.2 Requests at the DNS Root

In this section we perform a similar analysis as for §4.1, in which we look into DNS queries received at some of the Root DNS servers, namely A-, B-, C- and D-Root, and create a distribution of the number of queries received per source IP address (*i.e.*, per recursive).




Figure 5: Distribution of the number of queries for the DS record of `nl` received for each recursive. Dataset: A, B, C, D DTIL on 2017-04-12t11:00Z for 3 hours.

In this analysis we use data from the DITL (*Day In The Life*) dataset of 2017, available at DNS-OARC [10]. We look at all DNS queries received for the DS record of the domain `nl`, during the interval between 11am to 1pm on April 12, 2017 (UTC). Note that the DS record for `nl` has a TTL of 86400 seconds (24 hours). That is, we do not expect to see multiple queries for the DS record of `nl` within the 3-hour interval at the selected DNS Root servers.

Figure 5 shows the distribution of the number of queries sent by recursives to the DS record of `nl`. We see that the majority (around 80%) of recursives does send only one query within the 3-hour interval. However, all the Root letters we looked into have seen multiple queries from 20% of recursives. Note that the x-axis is truncated, and we see up to around 360 queries from a single recursive (*i.e.*, source IP address) within the 3-hour interval. (In this analysis, for the 3-hour interval of measured data, the sample of unique recursives varies from 982 for B-Root to 2094 recursives for C-Root.)

5 THE CLIENT’S VIEW OF AUTHORITATIVES UNDER DDOS

We next evaluate the effects of DDoS attacks on client experiments. This work builds our studies of caching in controlled experiments (§3) and passive observations (§4) have shown that it often works, but not always—about 70% of controlled experiments and 30% of passive observations see full cache lifetimes. We first consider *complete* failure of authoritative servers, then *partial* failure.

5.1 Emulating DDoS

To emulate DDoS attacks we begin with the same test domain (`cachetest.nl`) we used for controlled experiments in §3.2. We run a normal DNS for some time, querying from RIPE Atlas. After caches are warm, we then simulate a DDoS attack by dropping some fraction or all incoming DNS queries to each authoritative. (We drop traffic with Linux iptables.) After we begin dropping traffic, answers come either from caches

at recursives or, for partial attacks, from a lucky query that passes through.

This emulation of DDoS captures traffic loss that occurs in DDoS attack as router queues overflow. This emulation is not perfect, since we simulate loss at the last hop-router, but in real DDoS attacks packets are often lost on access links near the target. Our emulation approximates this effect with one aggregate loss rate.

DDoS attacks are also accompanied by queueing delay, since buffers at and near the target are full. We do not model queueing delay, although we do observe latency due to retries. In modern routers, queueing delay due to full router buffers should be less than the retry interval. In addition, observations during real-world DDoS events show that the few queries that are successful see response times that are not much higher than typical [22], suggesting that loss (and not delay) is the dominant effect of DDoS in practice.


5.2 Clients During Complete Authoritatives Failure

We first evaluate the worst-case scenario: complete unreachability of all authoritative name servers. Our goal is to understand when and for how long caches cover such an outage.


Table 4 shows the experiments we executed – being A, B, and C the ones which simulate complete failure. In Experiment A, each VP makes only one query before the DDoS begins. In Experiment B we allow several queries to take place, and Experiment C allows several queries with a shorter TTL.

Caches Protect Some: We first consider Experiment A, with one query that warms the cache immediately followed by the attack. Figure 6a shows these responses over time, with the onset of the attack the first downward arrow between 0 and 10 minutes, and with the cache expired after the second downward arrow between 60 and 70 minutes. We see that after the DDoS starts but before the cache has fully expired (between the downward arrows) initially 30% and eventually 65% of queries fail with either no answer or a SERVFAIL error message. While not good, this does mean that 35% to 70% of queries during the DDoS are successfully served from the cache. By contrast, shortly after the cache expires, almost all queries succeed (only 25 VPs or 0.2% of the total seem to provide stale answers).


Caches Fill at Different Times: In a more realistic scenario, VPs have filled their caches at different times. In Experiment A, caches are freshly filled and should last for a full hour after the start of attack. Experiment B is designed for the opposite and worst case: we begin warming the cache one hour before the attack and query 6 times of each VP. Other parameters are the same, with the attack lasting for 60 minutes (also the cache duration), but then we restore the authoritatives to service.



(a) Experiment A: 3600-10min-1down; arrows indicate DDoS start and cache expiration



(b) Experiment B: 3600-10min-1down-1up; arrows indicate DDoS start and recovery



(c) Experiment C: 1800-10min-1down-1up; arrows indicate DDoS start, cache expiration and recovery

Figure 6: Answers received during DDoS attacks.

Figure 6b shows the results of Experiment B. While about 50% of VPs are served from the cache in the first 10 minute round after the DDoS starts, the fraction served drops quickly and is at only about 3% one hour later. Three factors are in play here: most caches were filled 60 minutes before the attack and are timing out in the first round. While the timeout and query rounds are both 60 minutes apart, Atlas intentionally spreads queries out over 5 minutes, so we expect that some queries happen after 59 minutes and others 61 minutes.

Second, we know some large recursives have fragmented caches (§3.5), so we expect that some of the successes between times 70 and 110 minutes are due to caches that were

Experiment Parameters							Results						
	TTL in sec.	DDoS start	DDoS dur.	queries before	total dur.	probe interval	failure	Total probes	Valid probes	VPs	Queries	Total answers	Valid answers
A	3600	10	60	1	120	10	100% (both NSes)	A 9224	8727	15339	136423	76619	76181
B	3600	60	60	6	240	10	100% (both NSes)	B 9237	8827	15528	357102	293881	292564
C	1800	60	60	6	180	10	100% (both NSes)	C 9261	8847	15578	258695	199185	198197
D	1800	60	60	6	180	10	50% (one NS)	D 9139	8708	15332	286231	273716	272231
E	1800	60	60	6	180	10	50% (both NSes)	E 9153	8708	15320	285325	270179	268786
F	1800	60	60	6	180	10	75% (both NSes)	F 9141	8727	15325	278741	259009	257740
G	300	60	60	6	180	10	75% (both NSes)	G 9206	8771	15481	274755	249958	249042
H	1800	60	60	6	180	10	90% (both NSes)	H 9226	8778	15486	269030	242725	241569
I	60	60	60	6	180	10	90% (both NSes)	I 9224	8735	15388	253228	218831	217979

Table 4: DDoS emulation experiments [34]; DDoS start, durations and probe interval are given in minutes.




Figure 7: Timeseries of answers for Experiment B

filled between times 10 and 50 minutes. This can actually be seen in Figure 7, where we show a timeseries of the answers for Experiment B, where we see CC (correct cache responses) between times 60 and 90.

Third, we see an increase in the number of CA queries that are answered by the cache with expired TTLs (Figure 7). This increase is due to servers serving stale content.

Caches Eventually All Expire: Finally, we carry out a third emulation but with half the cache lifetime (1800 s or 30 minutes rather than the full hour). Figure 6c shows response over time. These results are similar to Experiment B, with rapid fall-off when the attack starts as caches age. After the attack has been underway for 30 minutes all caches must have expired and we see only a few (about 2.6%) residual successes.

5.3 Discussion of Complete Failures

Overall we see that *caching is partially successful in protecting during a DDoS*. With full, valid caches, half or more VPs get service. However, caches are filled at different times and expire, so an operator cannot count on a full cache duration for any customers, even for popular (“always in the cache”) domains. The protection provided by caches depends on their state in the recursive resolver, something outside the

operator’s control. In addition, our evaluation of caching in §3 showed that caches will end early for some VPs.

Second, we were surprised that a tiny fraction of VPs are successful after all caches should have timed out (after the 80 minutes period in Experiment A, and between 90 and 110 minutes in Experiment C). These successes suggest an early deployment of “serve stale”, something currently under review in the IETF [18] is to serve a previously known record beyond its TTL if authoritatives are unreachable, with the goal of improving resilience under DDoS. We investigated the Experiment C, where we see that 732 answers of the 2390 successes in the second half of the outage. These successes are from two 4 IP addresses corresponding to the OpenDNS and Google public DNS servers, suggesting experimentation not yet widespread.

5.4 Client Reliability During Partial Authoritative Failure

The previous section examined DDoS attacks that result in complete failure of all authoritatives, but often DDoS attacks result in *partial failure*, with 50% or 90% packet loss. (For example, consider the November 2015 DDoS attack on the DNS Root [22].) We next study experiments with partial failures, showing that *caching and retries together nearly fully protect 50% DDoS events, and protect half of VPs even during 90% events*.

We carry out several Experiments D to I in Table 4 We follow the procedure outlined in §5.1, looking at the DDoS-driven loss rates of 50%, 75%, and 90% with TTLs of 1800 s, 300 s and 60 s. Graphs omitted due to space can be found in Appendix D.

Near-Full Protection from Caches During Moderate Attacks: We first consider Experiment E, a “mild” DDoS with 50% loss, a with VP success over time in Figure 8a. In spite of a loss rate that would be crippling to TCP, nearly all VPs are successful in DNS. This success is due to two factors: first, we know that many clients are served from caches, as was shown

in Experiment A with full loss (Figure 6a). Second, most recursive retry queries, so they recover from loss of a single packet and are able to provide an answer. Together, these mean that failures during the first 30 minutes of the event is 8.5%, slightly higher than the 4.8% fraction of failures before the DDoS. For this experiment, the TTL is 1800 s (30 minutes), so we might expect failures to increase halfway through the DDoS. We do not see any increase in failures because caching and retries are *synergistic*, a successful retried query will place the answer in a cache for a later query. The importance of this result is that *DNS can survive moderate-size attacks when caching is possible*. While a positive, retries do increase latency, something we study in §5.5.

Attack Intensity Matters: While clients do quite well with 50% loss at all authoritatives, failures increase with the intensity of the attack.

Experiments F and H, shown in Figure 8b and Figure 8c increase the loss rate to 75% and 90%. We see the number of failures increases to about 19.0% with 75% loss and 40.3% with 90% loss. It is important to note that *roughly 60% the clients are still served even with 90% loss*.


We also see that this level of success is consistent over the entire hour-long DDoS event, even though the cache duration is only 30 minutes. This consistency confirms the importance of caching and retries in combination.

To verify the effects of this interaction, Experiment I changes the caching duration to 60 s, less than one round or probing. Comparing Experiment I in Figure 8d with H in Figure 8c, we see that the failure rate increases from 30% to about 63%. Still, with no caching in fact, we can see that roughly 37% of queries still are answered. We investigate retries in §6.


5.5 Client Latency During Partial Authoritative Failure

We showed that client reliability is higher than expected during failures (§5.4) due to a combination of caching and retries. We next consider client *latency*. Latency will increase during the DDoS because of retries and queuing delay, but we will show that latency increases less than one might expect due to caching.


To examine latency we return to Experiments D through I (Table 4), but look at latency (time to complete a query) rather than success. Another metric to measure client’s experience during a DDoS is how long it takes to resolve queries, *i.e.*, the round-trip time (RTT). Figures 9a to 9d show latency during each emulated DDoS scenario (experiments with figures omitted here are in Appendix D Latencies not evenly distributed, since some requests get through immediately while others must be retried one or more times, so in addition to mean, we show 50, 75 and 90% quantiles to characterize the tail of the distribution.




(a) Experiment E (1800-50p-10min): 50% packet loss



(b) Experiment F (1800-75p-10min): 75% packet loss




(c) Experiment H (1800-90p-10min): 90% packet loss




(d) Experiment I (60-90p-10min): 90% packet loss


Figure 8: Answers received during DDoS attacks; 1st and 2nd vertical lines show start and end of DDoS.




(a) Experiment E: 50% packet loss (1800s TTL)



(b) Experiment F: 75% packet loss (1800s TTL)



(c) Experiment H: 90% packet loss (1800s TTL)



(d) Experiment I: 90% packet loss (60s TTL)

Figure 9: Latency results; Shaded area indicates the interval of an ongoing DDoS attack.

We emulate DDoS by dropping requests (§5.1) and, hence, latencies reflect retries and loss, but not queueing delay, underrepresenting latency in real-world attacks. However, their shape (some low latency and a few long) is consistent with and helps explain what has been seen in the past [22].

Beginning with Experiment E, the moderate attack in Figure 9a, we see *no* change to median latency. This result is consistent with many queries being handled by the cache, and half of those not handled by the cache getting through anyway. We do see higher latency in the 90%ile tail, reflecting successful retries. This tail also increases the mean some.

This trend increases in Experiment F in Figure 9b, where 75% of queries are lost. Now we see the 75%ile tail has increased, as has the number of unanswered queries, and the 90%ile is twice as long as in Experiment E.

We see the same latency in Experiment H with DDoS causing 90% loss. RIPE Atlas has a fixed number of timeouts, so the larger attack results in more unsuccessful queries, but latency for successful queries is not much worse than with 75% loss. Median latency is still low due to cached replies.

Finally, Experiment I greatly reduces opportunities for caching by reducing cache lifetime to one minute. Figure 9d shows that loss of caching increases median RTT and significantly increases the tail latency. Compared with Figure 9c (same packet loss ratio but 1800 s TTL), we can clearly see the benefits of caching in terms of latency (in addition to reliability): a half-hour TTL value reduced the latency from 1300 ms to 390 ms.


Summary: DDoS effects often increase client latency. For moderate attacks, increased latency is seen only by a few “unlucky” clients whose do not see a full cache and whose queries are lost. Caching has an important role in reducing latency during DDoS, but while it can often mitigate most reliability problems, it cannot avoid latency penalties for all VPs. Even when caching is not available, roughly 40% of clients get an answer, either by serving stale or retries as we investigate next.

6 THE AUTHORITATIVE’S PERSPECTIVE


Results of partial DDoS events (§5.4) show that DNS is surprisingly reliable well—even with a DDoS resulting in 90% packet loss, about 40% of VPs get answers (Figure 8d). We suggested this success is possible because of a combination of caching and retries. We next examine this from the perspective of the authoritative server.

6.1 Recursive-Authoritative Traffic during a DDoS


We first ask: what is the cost of these retries on the authoritative servers? To investigate this question, we return the partial DDoS experiments and look at how many queries are sent to the server. We measure queries before they are



(a) Experiment F: 1800-75p-10min, 75% packet loss



(b) Experiment H: 1800-90p-10min, 90% packet loss



(c) Experiment I: 60-90p-10min, 90% packet loss

Figure 10: Number of received queries by the authoritative servers. Shaded area indicates the interval of an ongoing DDoS attack.

dropped by our simulated DDoS. Recursives must make multiple queries to resolve a name. We break out each type of query: for the nameserver (NS), the nameserver’s IPv4 and v6 addresses (A-for-NS and AAAA-for-NS), and finally the desired query (AAAA-for-PID). Note that the authoritative is IPv4 only; so AAAA-for-NS is non-existent and subject to negative caching, while the other records exist and use regular caching.

We begin with the DDoS causing 75% loss in Figure 10a. During the DDoS, queries increase by about 3.5 \times . We expect 4 trials, since the expected number of tries until success with loss rate p is $(1-p)^{-1}$. For this scenario, results are cached for

up to 30 minutes, so successful queries are reused in recursive caches. This increase occurs both for the target AAAA record, and also for the non-existent AAAA-for-NS records. Negative caching for our zone is configured to 60 s, making caching of NXDOMAINs for AAAA-for-NS less effective than positive caches.

The offered load on the server increases further with more loss (90%), as shown in Experiment H (Figure 10b). The higher loss rate results in a much higher offered load on the server, average 8.2 \times normal.

Finally, in Figure 10c we reduce the effects of caching at a 90% DDoS and with a TTL of 60s. Here we see also about 8.1 \times more queries at the server before the attack. Comparing this case to Experiment H, caching reduces the offered load on the server by about 40%.

Implications: The implication of this analysis is that legitimate clients “hammer” with retries the already-stressed server during a DDoS. For clients, retries are important to get reliability; and each client independently chooses to retry.

The server is already under stress due to the DDoS, so these retries add to that stress. However, a server experiencing 90% loss is already at 10 \times its capacity, and maximum capacity is often 10 \times normal load (for example, see [7]), so additional traffic of 10 \times normal load is perhaps increases the attack traffic by only 10%. This multiplier depends on the implementations stub and recursive resolver, as well as application-level retries and deflection (users hitting reload in their browser, and later giving up). Our experiment omits application-level retries and likely gives a lower bound. We next examine specific recursive implementations to see their behavior.

6.2 Sources of Retries: Software and Multi-level Recursives

Experiments in the prior section showed that recursive resolvers “hammer” authoritatives when queries are dropped. We reexamine DNS software (since 2012 [50]), and additionally show deployments amplify retries.

Recursive Software: Prior work showed that recursive servers retry many times when an authoritative is unresponsive [50], with evaluation of BIND 9.7 and 9.8, DNSCache, Unbound, WindowsDNS and PowerDNS. We studied retries in BIND 9.10.3 and Unbound 1.5.8 to quantify the number of retries. Examining only requests for AAAA records, we see that normal requests with a responsive authoritative ask for the AAAA records for all authoritatives and the target name (3 total requests when there are 2 authoritatives). When all authoritatives are unavailable, we see about 7 \times more requests before the recursives time out. (Exact numbers vary in different runs, but typically each request is made 6 or 7 times.) Such retries are appropriate, provided they are paced (both use exponential backoff), they explain part of




Figure 11: *Rn* recursives and AAAA queries used in Experiment I, normalized by the number of probe IDs.

the increase in legitimate traffic during DDoS events. Full data is in [Appendix G](#).

Recursive Deployment: Another source of extra retries is complex recursive deployments. We showed that operators of large recursives often use complex, multi-level resolution infrastructure (§3.5). This infrastructure can amplify the number of retries during reachability problems at authoritatives.

To quantify amplification, we count both the number of *Rn* recursives and AAAA queries for each probe ID reaching our authoritatives. [Figure 11](#) show the results. These values represent the amplification in two ways: during stress, more *Rn* recursives will be used for each probe ID and these *Rn* will generate more queries to the already stressed authoritatives. As the figures shows, the median number of *Rn* recursives employed doubles (from 1 to 2) during the DDoS event, as does the 90%ile (from 2 to 4). The maximum rises to 39. The number of queries for each probe ID grows more than 3 \times , from 2 to 7. Worse, the 90%ile grows more than 6 \times (3 queries to 18). The maximum grows 53.5 \times , reaching up to 286 queries for one single probe ID. This value, however, is a lower bound, given there is large number of A and AAAA queries that ask for NS records, and not the probe ID (AAAA and A-for NS in [Figure 10](#)).

Most complex resolution infrastructure is proprietary (as far as we know only one study has examined them [43]), so we cannot make recommendations about how large recursive resolvers ought to behave. We suggest that the aggregate traffic of large resolvers should strive to be within a constant factor of single resolvers, perhaps a factor of 4. We also encourage additional study of large resolvers, and their operators to share information about their behavior.

7 RELATED WORK

Caching by Recursives: Several groups have shown that DNS caching can be imperfect. Hao and Wang analyzed the impact of nonce domains on DNS recursive’s caches [13]. Using two weeks of data from two universities they showed that filtering one-time domains improves cache hit rates. In

two studies, Pang *et al.* [27, 28] reported that web clients and local recursives do not always honor TTL values provided by authoritatives. Almeida *et al.* [2] analyzed DNS traces of a mobile operator, and used a mobile application to see TTLs in practice. They find that most domains have short TTLs (less than 60 s), and report and evidence of TTL manipulation by recursives. Schomp *et al.* [43] demonstrate widespread use of multi-level recursives by large operators, as well as TTL manipulation. Our work builds on this prior work, examining caching and TTL manipulation systematically and considering its effects on resilience.

DNS client behavior: Yu *et al.* investigated how stubs and recursives select authoritative servers, and were the first to demonstrate the large number of retries when all authoritatives are unavailable [50]. We also [24] investigated how recursives select authoritative servers in the wild and found that recursives tend to prefer authoritatives with shorter latency, but query all authoritatives for diversity. We confirm Yu’s work and focus on authoritative selection during DDoS from several perspectives.

Authoritatives during DDoS: We investigated how the Root DNS service behaved during the Nov. 2015 DDoS attacks [22]. This report focuses on the interactions of IP anycast and both latency and reachability, as seen from RIPE Atlas. Rather than look at aggregate behavior and anycast, our methodology here examines how clients interact with their recursive resolvers, while this prior work focused on authoritatives only, bypassing recursives. In addition, here we have full access to clients and authoritatives traffic during our experiments, and we evaluate DDoS with controlled loss rates. The prior study has incomplete data and focuses on specific results of two events. These differences stem from their study of natural experiments from real-world events and our controlled experiments.

8 IMPLICATIONS

We evaluated DNS resilience, showing that caches and retries can cover up the effects of a DDoS attack; provided the cache is full and some requests can get to authoritative servers.

Recent attacks on DNS services have had somewhat different outcomes for users. The Root Server System has often been a target, with recent attacks in Nov. 2015 [37] and June 2016 [38]. The DNS Root has 13 authoritative “letters”, each an authoritative “server” implemented with some or many IP anycast instances. Analysis of these DDoS events showed that their effects were uneven across letters, with some or all anycast instances of many letters showing high loss, while others showed little or no effects (see Moura *et al.* [22] for details). However, the Root Operators report “There are no known reports of end-user visible error conditions during, and as a result of, this incident. Because the DNS protocol is designed to cope with partial reachability...” [37].

In Oct. 2016, a much larger attack was directed at Dyn, a provider of DNS service for many second-level domains [14]. Although Dyn has a capable infrastructure and immediately took steps to address service problems, there were reports of user-visible service disruption in the technical and even popular press [30]. Reports describe intermittent failure of prominent websites including “Twitter, Netflix, Spotify, Airbnb, Reddit, Etsy, SoundCloud and The New York Times”.

Our work can partially explain these very different outcomes. The Root DNS saw few or no user-visible problems because data in the root zone is cachable for a day or more, and because multiple letters and many anycast instances were continuously available. (All measurements in this paragraph are as of 2018-05-22.) Records in the root zone have TTLs of 1 to 6 days, and www.root-servers.org reports 922 anycast instances operating across the 13 authoritative servers. Dyn also operates a large infrastructure (<https://dyn.com/dns/network-map/> reports 20 “facilities”), and faced a larger attack (reports of 1.2 Tb/s [42], compared to estimates of 35 Gb/s for the Nov. 2015 root attack [22]). But a key difference is *all* of the Dyn’s customers listed above use DNS-based CDNs (for a description, see [8]) with multiple, Dyn-hosted DNS components with TTLs that range from 120 to 300 s.

Our experiments can explain the root cause behind these different outcomes. Users of the root benefited from caching and saw performance like Experiment E (Figure 8a), because root contents (TLDs like .com and country codes) are popular and certainly cached in recursives, and because some root letters were always available to refresh caches. By contrast, users requiring domains with very short TTLs (like the websites that had problems) receive performance more like Experiment I (Figure 8d) or Experiment C (Figure 6c). Even though these services are very popular, short TTLs allow them very little benefit from caching during an extended interruption of DNS service.

This example shows the importance DNS’s multiple methods of resilience (caching, retries, and at least partial availability of one authoritative). It suggests that CDN operators may wish to consider longer timeouts (Experiment H suggests 30 minutes, Figure 8c), to allow caching to serve some role, giving DNS operators deploy defenses.

Finally, this evaluation helps complete our picture of DNS latency and reliability for DNS services that may consist of multiple authoritatives, some or all using IP anycast with multiple sites. To minimize latency, prior work has shown a single authoritative using IP anycast should maximize geographic dispersion of sites [41]. The latency of an overall DNS service with *multiple* authoritatives can be limited by the one with largest latency [24]. Prior work about resilience to DDoS attack has shown that individual IP anycast sites will suffer under DDoS as a function of the attack traffic that site receives relative to its capacity [22]. We show that

the overall reliance of a DNS service composed of multiple authoritatives using IP anycast tends to be as resilient as the *strongest* individual authoritative. The reason for these opposite results is that, in both cases, recursive resolvers will try *all* authoritatives of a given service. For latency, they will sometimes chose a distant authoritative, but for resilience, they will continue until they find the most available authoritative.

9 CONCLUSIONS

This paper represents the first study of how the DNS resolution system behaves when authoritative servers are under DDoS attack. Caching and retries at recursive resolvers are key factors in this behavior. We show that together, caching and retries by recursive resolvers greatly improve the resilience of the DNS as a whole. In fact, they can largely cover over partial DDoS attacks for many users. As one example, we show that half of users continue to receive service even when all authoritative servers for a given domain are under a DDoS that causes 90% request loss. The primary cost of DDoS for users can be greater latency, but even this penalty is uneven across users, with a few getting much greater latency while some see no or little change. Finally, we show that one result retries is that traffic from legitimate users to authoritatives greatly increases (up to 8×) during service interruption, and that this effect is magnified by complex, multi-layer recursive resolver systems. The key outcome of work is to quantify the importance of caching in recursives to resilience, encouraging use of at least moderate TTLs wherever possible.

Acknowledgments

The authors would like to thank Jelte Jansen, Benno Overeinder, Marc Groeneweg, and Wes Hardaker for their valuable comments on paper drafts.

Giovane C. M. Moura and Moritz Müller, and Marco Davids developed this work as part of the SAND project (<http://www.sand-project.nl>).


This research has been partially supported by measurements obtained from RIPE Atlas, an open measurements platform operated by RIPE NCC.

John Heidemann’s research is partially sponsored by the Air Force Research Laboratory and the Department of Homeland Security under agreements number FA8750-17-2-0280 and FA8750-17-2-0096. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon.


REFERENCES

- [1] 1.1.1.1. The Internet’s Fastest, Privacy-First DNS Resolver. <https://1.1.1.1/>, Apr. 2018.
- [2] ALMEIDA, M., FINAMORE, A., PERINO, D., VALLINA-RODRIGUEZ, N., AND VARVELLO, M. Dissecting DNS Stakeholders in Mobile Networks. In *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies* (New York, NY, USA, 2017), CoNEXT


- '17, ACM, pp. 28–34.
- [3] ANTONAKAKIS, M., APRIL, T., BAILEY, M., BERNHARD, M., BURSSTEIN, E., COCHRAN, J., DURUMERIC, Z., HALDERMAN, J. A., INVERNIZZI, L., KALLITSIS, M., KUMAR, D., LEVER, C., MA, Z., MASON, J., MENSCHER, D., SEAMAN, C., SULLIVAN, N., THOMAS, K., AND ZHOU, Y. Understanding the Mirai botnet. In *Proceedings of the 26th USENIX Security Symposium* (Vancouver, BC, Canada, Aug. 2017), USENIX, pp. 1093–1110.
 - [4] ARBOR NETWORKS. Worldwide infrastructure security report. Tech. Rep. 2012 Volume VIII, Arbor Networks, Sept. 2012.
 - [5] BAJPAI, V., ERAVUCHIRA, S., SCHÖNWÄLDER, J., KISTELEKI, R., AND ABEN, E. Vantage Point Selection for IPv6 Measurements: Benefits and Limitations of RIPE Atlas Tags. In *IFIP/IEEE International Symposium on Integrated Network Management (IM 2017)* (Lisbon, Portugal, May 2017).
 - [6] BAJPAI, V., ERAVUCHIRA, S. J., AND SCHÖNWÄLDER, J. Lessons learned from using the RIPE Atlas platform for measurement research. *SIGCOMM Comput. Commun. Rev.* 45, 3 (July 2015), 35–42.
 - [7] BUSH, R., KARRENBERG, D., KOSTERS, M., AND PLZAK, R. Root name server operational requirements. RFC 2870, Internet Request For Comments, June 2000. (also Internet BCP-40).
 - [8] CALDER, M., FLAVEL, A., KATZ-BASSETT, E., MAHAJAN, R., AND PADHYE, J. Analyzing the performance of an anycast CDN. In *Proceedings of the ACM Internet Measurement Conference* (Tokyo, Japan, Oct. 2015), ACM.
 - [9] DE OLIVEIRA SCHMIDT, R., HEIDEMANN, J., AND KUIPERS, J. H. Anycast Latency: How Many Sites Are Enough? In *Passive and Active Measurements (PAM)* (2017), pp. 188–200.
 - [10] DNS OARC. DITL Traces and Analysis. <https://www.dns-oarc.net/index.php/oarc/data/ditl/2018>, Apr. 2018.
 - [11] ELZ, R., BUSH, R., BRADNER, S., AND PATTON, M. Selection and Operation of Secondary DNS Servers. RFC 2182 (Best Current Practice), July 1997.
 - [12] GOOGLE. Public DNS. <https://developers.google.com/speed/public-dns/>, Jan. 2018.
 - [13] HAO, S., AND WANG, H. Exploring Domain Name Based Features on the Effectiveness of DNS Caching. *SIGCOMM Comput. Commun. Rev.* 47, 1 (Jan. 2017), 36–42.
 - [14] 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.
 - [15] ICANN. RSSAC002: RSSAC Advisory on Measurements of the Root Server System. <https://www.icann.org/en/system/files/files/rssac-002-measurements-root-20nov14-en.pdf>, Nov. 2014.
 - [16] KOTTLER, S. February 28th DDoS Incident Report | Github Engineering, Mar. 2018. <https://githubengineering.com/ddos-incident-report/>.
 - [17] LAWRENCE, D., AND KUMARI, W. Serving Stale Data to Improve DNS Resiliency. <https://tools.ietf.org/html/draft-tale-dnsop-serve-stale-02>, Oct. 2017.
 - [18] LAWRENCE, D., AND KUMARI, W. Serving Stale Data to Improve DNS Resiliency-02. Internet Draft, Oct. 2017. <https://www.ietf.org/archive/id/draft-tale-dnsop-serve-stale-02.txt>.
 - [19] MOCKAPETRIS, P. Domain names - concepts and facilities. RFC 1034 (Internet Standard), Nov. 1987. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 2065, 2181, 2308, 2535, 4033, 4034, 4035, 4343, 4035, 4592, 5936, 8020.
 - [20] MOCKAPETRIS, P. Domain names - implementation and specification. RFC 1035 (Internet Standard), Nov. 1987. Updated by RFCs 1101, 1183, 1348, 1876, 1982, 1995, 1996, 2065, 2136, 2181, 2137, 2308, 2535, 2673, 2845, 3425, 3658, 4033, 4034, 4035, 4343, 5936, 5966, 6604, 7766.
 - [21] MORALES, C. February 28th DDoS Incident Report | Github Engineering NETSCOUT Arbor Confirms 1.7 Tbps DDoS Attack; The Terabit Attack Era Is Upon Us, Mar. 2018. <https://www.arbornetworks.com/blog/asert/netscout-arbor-confirms-1-7-tbps-ddos-attack-terabit-attack-era-upon-us/>.
 - [22] MOURA, G. C. M., DE O. SCHMIDT, R., HEIDEMANN, J., DE VRIES, W. B., MÜLLER, M., WEI, L., AND HESSELMAN, C. Anycast vs. DDoS: Evaluating the November 2015 root DNS event. In *Proceedings of the ACM Internet Measurement Conference* (Nov. 2016).
 - [23] MOURA, G. C. M., HEIDEMANN, J., MÜLLER, M., DE O. SCHMIDT, R., AND DAVIDS, M. Datasets from “when the dike breaks: Dissecting DNS defenses during DDoS”. Web page https://ant.isi.edu/datasets/dns/#Moura18a_data, May 2018.
 - [24] MÜLLER, M., MOURA, G. C. M., DE O. SCHMIDT, R., AND HEIDEMANN, J. Recursions in the wild: Engineering authoritative DNS servers. In *Proceedings of the ACM Internet Measurement Conference* (London, UK, 2017), pp. 489–495.
 - [25] OPENDNS. Setup Guide: OpenDNS. <https://www.opendns.com/setupguide/>, Jan. 2018.
 - [26] PAN, J., HOU, Y. T., AND LI, B. An overview of dns-based server selections in content distribution networks. *Computer Networks* 43, 6 (2003), 695–711.
 - [27] PANG, J., AKELLA, A., SHAIKH, A., KRISHNAMURTHY, B., AND SESHAN, S. On the Responsiveness of DNS-based Network Control. In *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement* (New York, NY, USA, 2004), IMC '04, ACM, pp. 21–26.
 - [28] PANG, J., HENDRICKS, J., AKELLA, A., DE PRISCO, R., MAGGS, B., AND SESHAN, S. Availability, usage, and deployment characteristics of the domain name system. In *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement* (New York, NY, USA, 2004), IMC '04, ACM, pp. 1–14.
 - [29] PAUL VIXIE AND GERRY SNEERINGER AND MARK SCHLEIFER. Events of 21-oct-2002, Oct. 2002. <http://c.root-servers.org/october21.txt>.
 - [30] PERLROTH, N. Hackers used new weapons to disrupt major websites across U.S. *New York Times* (Oct. 22 2016), A1.
 - [31] PERLROTH, N. Tally of cyber extortion attacks on tech companies grows. *New York Times Bits Blog*, <http://bits.blogs.nytimes.com/2014/06/19/tally-of-cyber-extortion-attacks-on-tech-companies-grows/>, June 2016.
 - [32] PETERSON, A. Ec2 resolver changing ttl on dns answers? Post on the DNS-OARC dns-operations mailing list, <https://lists.dns-oarc.net/pipermail/dns-operations/2017-November/017043.html>, Nov. 2017.
 - [33] QUAD9. Quad9 | Internet Security & Privacy In a Few Easy Steps. <https://quad9.net>, Jan. 2018.
 - [34] RIPE NCC. RIPE Atlas measurement IDS. <https://atlas.ripe.net/measurements/ID>, Dec. 2017. ID is the experiment ID: TTL60: 10443671, TTL1800: 10507676, TTL3600: 10536725, TTL86400: 10579327, TTL3600-10min: 10581463, A:10859822, B: 11102436, C :11221270, D:11804500, E: 11831403, F: 11831403, G: 12131707, H:12177478, I: 12209843.
 - [35] RIPE NCC STAFF. RIPE Atlas: A Global Internet Measurement Network. *Internet Protocol Journal (IPJ)* 18, 3 (Sep 2015), 2–26.
 - [36] RIPE NETWORK COORDINATION CENTRE. RIPE Atlas - Raw data structure documentations, https://atlas.ripe.net/docs/data_struct/, 2018.
 - [37] ROOT SERVER OPERATORS. Events of 2015-11-30, Nov. 2015. <http://root-servers.org/news/events-of-20151130.txt>.
 - [38] ROOT SERVER OPERATORS. Events of 2016-06-25. Tech. rep., Root Server Operators, June 29 2016.
 - [39] ROOT SERVER OPERATORS. Root DNS, Feb. 2017. <http://root-servers.org/>.
 - [40] SANTANNA, J. J., VAN RIJSWIJK-DEIJ, R., HOFSTEDE, R., SPEROTTO, A., WIERBOSCH, M., GRANVILLE, L. Z., AND PRAS, A. Booters—an analysis of DDoS-as-a-Service attacks. In *Proceedings of the 14th IFIP/IEEE International Symposium on Integrated Network Management* (Ottawa, Canada, May 2015), IFIP.
 - [41] SCHMIDT, R. D. O., HEIDEMANN, J., AND KUIPERS, J. H. Anycast latency:




(a) TTL 60 s




(b) TTL 1800 s



(c) TTL 3600 s



(d) TTL 86400 s (1 day)



(e) TTL 3600 s (1 hour), Probing Interval T=10min

Figure 12: Fraction of query answer types over time

How many sites are enough? In *Proceedings of the Passive and Active Measurement Workshop* (Sydney, Australia, Mar. 2017), Springer,

- pp. 188–200.
- [42] SCHNEIER, B. Lessons from the Dyn DDoS attack. blog https://www.schneier.com/essays/archives/2016/11/lessons_from_the_dyn.html, Nov. 2016.
- [43] SCHOMP, K., CALLAHAN, T., RABINOVICH, M., AND ALLMAN, M. On measuring the client-side DNS infrastructure. In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference* (Oct. 2013), ACM, pp. 77–90.
- [44] SENGUPTA, S. After threats, no signs of attack by hackers. *New York Times* (Apr. 1 2012), A1.
- [45] SIDN LABS. .nl stats and data, Mar. 2017. <http://stats.sidnlabs.nl/#network>.
- [46] UNBOUND. Unbound Documentation. <https://www.unbound.net/documentation/unbound.conf.html>, Jan. 2018.
- [47] WEI, L., AND HEIDEMANN, J. Does anycast hang up on you? In *IEEE International Workshop on Traffic Monitoring and Analysis (Dublin, Ireland)* (June 2017).
- [48] 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.
- [49] WULLINK, M., MOURA, G. C., MÜLLER, M., AND HESSELMAN, C. Entrada: A high-performance network traffic data streaming warehouse. In *Network Operations and Management Symposium (NOMS), 2016 IEEE/IFIP* (Apr. 2016), IEEE, pp. 913–918.
- [50] YU, Y., WESSELS, D., LARSON, M., AND ZHANG, L. Authority server selection in dns caching resolvers. *SIGCOMM Comput. Commun. Rev.* 42, 2 (Mar. 2012), 80–86.

A TTL MANIPULATIONS ACROSS DATASETS

In §3.4, we saw that there are a number of AC-type answers—queries that are answered by the authoritative even though we expected them to be served from a cache. We now investigate the relation between the number of AC-answers with regards the TTL used for the DNS records.

Figure 12 shows response types over time for our four different TTLs. Analyzing these figures, we can see that with exception of a 60 s, where all queries go to the AA, the number of AC answers is relatively constant. This consistency suggests cache fragmentation across the datasets. We also see that AA and CC values alternate, given as caches entries expires new queries are forwarded to the authoritatives.

B LIST OF PUBLIC RESOLVERS

We use the following list of public resolvers in §3.4. This list is obtained by searching DuckDuckGo for “public” and “dns” on 2018-01-06.

198.101.242.72	Alternate DNS
23.253.163.53	Alternate DNS
205.204.88.60	BlockAid Public DNS (or PeerDNS)
178.21.23.150	BlockAid Public DNS (or PeerDNS)
91.239.100.100	Censurfridns
89.233.43.71	Censurfridns
2001:67c:28a4::	Censurfridns
2002:d596:2a92:1:71:53::	Censurfridns
213.73.91.35	Chaos Computer Club Berlin

209.59.210.167	Christoph HochstÄdtter	66.244.95.20	OpenNIC
85.214.117.11	Christoph HochstÄdtter	207.192.69.155	OpenNIC
212.82.225.7	ClaraNet	72.14.189.120	OpenNIC
212.82.226.212	ClaraNet	2001:470:8388:2:20e:2eff:fe63:d4a9	OpenNIC
8.26.56.26	Comodo Secure DNS	2001:470:1f07:38b::1	OpenNIC
8.20.247.20	Comodo Secure DNS	2001:470:1f10:c6::2001	OpenNIC
84.200.69.80	DNS.Watch	194.145.226.26	PowerNS
84.200.70.40	DNS.Watch	77.220.232.44	PowerNS
2001:1608:10:25::1c04:b12f	DNS.Watch	9.9.9.9	Quad9
2001:1608:10:25::9249:d69b	DNS.Watch	2620:fe::fe	Quad9
104.236.210.29	DNSReactor	195.46.39.39	SafeDNS
45.55.155.25	DNSReactor	195.46.39.40	SafeDNS
216.146.35.35	Dyn	193.58.251.251	SkyDNS
216.146.36.36	Dyn	208.76.50.50	SmartViper Public DNS
80.67.169.12	FDN	208.76.51.51	SmartViper Public DNS
2001:910:800::12	FDN	78.46.89.147	ValidDOM
85.214.73.63	FoeBud	88.198.75.145	ValidDOM
87.118.111.215	FoolDNS	64.6.64.6	Verisign
213.187.11.62	FoolDNS	64.6.65.6	Verisign
37.235.1.174	FreeDNS	2620:74:1b::1:1	Verisign
37.235.1.177	FreeDNS	2620:74:1c::2:2	Verisign
80.80.80.80	Freenom World	77.109.148.136	Xiala.net
80.80.81.81	Freenom World	77.109.148.137	Xiala.net
87.118.100.175	German Privacy Foundation e.V.	2001:1620:2078:136::	Xiala.net
94.75.228.29	German Privacy Foundation e.V.	2001:1620:2078:137::	Xiala.net
85.25.251.254	German Privacy Foundation e.V.	77.88.8.88	Yandex.DNS
62.141.58.13	German Privacy Foundation e.V.	77.88.8.2	Yandex.DNS
8.8.8.8	Google Public DNS	2a02:6b8::feed:bad	Yandex.DNS
8.8.4.4	Google Public DNS	2a02:6b8:0:1::feed:bad	Yandex.DNS
2001:4860:4860::8888	Google Public DNS	109.69.8.51	puntCAT
2001:4860:4860::8844	Google Public DNS		
81.218.119.11	GreenTeamDNS		
209.88.198.133	GreenTeamDNS		
74.82.42.42	Hurricane Electric		
2001:470:20::2	Hurricane Electric		
209.244.0.3	Level3		
209.244.0.4	Level3		
156.154.70.1	Neustar DNS Advantage		
156.154.71.1	Neustar DNS Advantage		
5.45.96.220	New Nations		
185.82.22.133	New Nations		
198.153.192.1	Norton DNS		
198.153.194.1	Norton DNS		
208.67.222.222	OpenDNS		
208.67.220.220	OpenDNS		
2620:0:ccc::2	OpenDNS		
2620:0:ccd::2	OpenDNS		
58.6.115.42	OpenNIC		
58.6.115.43	OpenNIC		
119.31.230.42	OpenNIC		
200.252.98.162	OpenNIC		
217.79.186.148	OpenNIC		
81.89.98.6	OpenNIC		
78.159.101.37	OpenNIC		
203.167.220.153	OpenNIC		
82.229.244.191	OpenNIC		
216.87.84.211	OpenNIC		

C CACHING OF REFERRALS


In §4.1, we analyzed the caching of A records for `ns[1-5].dns.nl`. In this section, we analyze the caching of *referrals* of two popular domains under `.nl`: `google.nl` and `www.google.nl`.

Whenever a recursive asks for the A record of `google.nl` to a `.nl` authoritative server, it does not get an IPv4 address as answer (or its A or AAAA record). Instead, it obtains NS records for the authoritative servers of `google.nl` (`dig @ns1.dns.nl A google.nl` illustrates that). In DNS, these NS answers to A queries are known as *referral*, since a client is referred to another authoritative server.


In this section, we analyze queries for A records of `google.nl` and `www.google.nl` – which are both answered with referrals by our `.nl` servers – using the same methodology as in §4.1. These referrals have a TTL of 3600 s (1 hour) and contain the NS records `ns1-ns4.google.com`.

After receiving these referrals, the client (or its recursive) will need to perform two tasks:

- (1) Retrieve A (or AAAA) records for `ns1-ns4.google.com`, which have a TTL of 4 days (172800 s).
- (2) Ask directly the IP address of the previous answer (#1) for the A record of `google.nl` and `www.google.nl`,



(a) Experiment D: 1800-50p-10min-ns1; arrows indicate DDoS start and recovery



(b) Experiment G: 300-75p-10min-2Ns; arrows indicate DDoS start and recovery

Figure 14: Answers received during DDoS attacks (extra graphs).




Figure 13: ECDF of Δt (ms) for recursives with at least 5 queries to the delegated domains google.nl and www.google.nl. TTL A record 300 s, TTL NS record 4 days, TTL delegation 1 hour.

which have a TTL of 300 s (5 minutes) at the time of this analysis.

Thus, the A record answer for google.nl and www.google.nl is valid for 300s. Once expired, the recursive will need to ask again to ns1-ns4.google.com. As such, if caching works

as expected, we should not see the same recursive asking for the A record answer for google.nl and www.google.nl to the authoritative servers of [.nl](http://nl) within one hour interval (given the answer will be referrals that are valid for 1 hour).

However, we observed at the authoritatives of [.nl](http://nl) more than 6,600 recursives that send 80,492 queries the A record of google.nl and www.google.nl before 1 hour, *i.e.*, the TTL of the referral.

Figure 13 shows the ECDF of time in between queries for the same recursive (Δt). By definition, in this measurement interval of 6 hours, we should expect that a recursive would send only up to 6 to the [.nl](http://nl) authoritatives. As such, any recursive that sends more than this already suggest that caching is not working as expected. Thus, in this figure, we only show recursives that are not caching as expected.

Analyzing Figure 13, we observe that 33% of the recursives re-send a query within 300 s which is \sim to the TTL of the A record google.nl and www.google.nl.

The implication of these findings for the operators of zones that contain many delegations, like TLD operators, is that their query load is not only influenced by the TTL of their records, but also by the TTL of their delegated zones (since caching does not work as expected all the time). The impact of a DDoS attack against authoritatives of delegating zones like [.nl](http://nl) becomes visible to many recursives already after the TTL of the delegated zones, like google.nl expires.

D EXTRA GRAPHS FROM PARTIAL DDOS

Although we carried out all experiments listed in Table 4 (§5), the body of the paper provides graphs for key results only. In this section, we include the graphs for experiments D and G.

Figure 14 shows the number of answers for experiments D and G, while Figure 15 shows the latency results. We can see at Figure 14a that when 1 NS fails (50%), there is no significant changes in the number of answered queries. Besides, most of the users will not notice in terms of latency (Figure 15a).

Figure 14b shows the results for Experiment G, in which the TTL is 300s (thus half of the probing interval). Similarly to experiment I, this experiment should not have caching active given the TTL is shorter. Moreover, it has a packet loss rate of 75%. At this rate, we can see that the far majority (72%) of users obtain an answer, with a certain latency increase (Figure 15b).

E VPS EXPERIENCE DURING FAILURE

In §5.3 we discuss complete failures. So far we have analyze the behavior of the entire measurement. However, a closer understanding of user experience can be understood by analyzing what each VP experiences. Table 5 shows the distribution of successful queries (correctly answered) for

	Min.	q1	Med.	Mean	q3	Max.	Max-AT	Samples
A	1	3	6	4.97	7	14	1	18
B	1	18	19	18.84	19	25	18	24
C	1	12	13	12.72	13	18	12	18

Table 5: Distribution of OK queries per VP (Max-AT refers to the number of queries that could be answered by authoritatives when authoritatives were (partially) reachable)

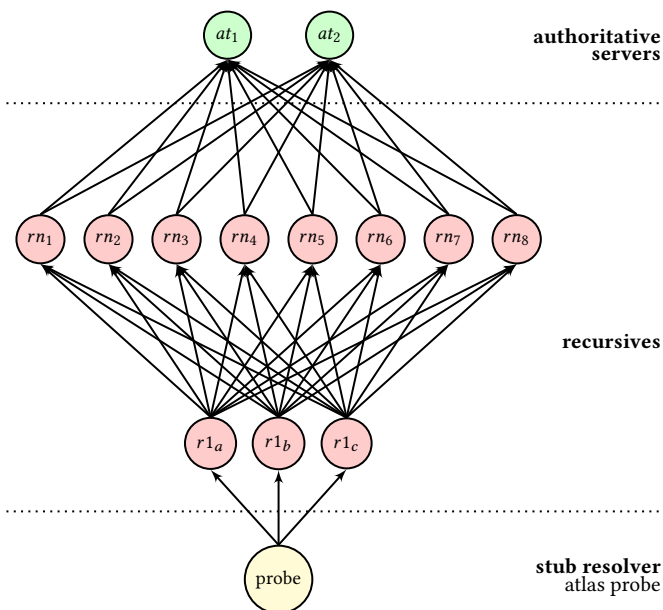




Figure 16: probe 284777 on dataset i. all $r1$'s are connected to all rn 's, which are connected to all at 's.



(a) Experiment D: 50% packet loss on 1 NS only (1800s TTL)



(b) Experiment G: 75% packet loss (300s TTL)

Figure 15: Latency results; Shaded area indicates the interval of an ongoing DDoS attack (extra graphs)

our datasets. The Max-authoritative value is the reference value we use for comparison: it denotes that maximum value of answers that could be achieved with no caching, only answered by the authoritatives (essentially, the number of measurement points minus the measurement points when authoritatives were totally unreachable). And samples value denotes the number of data points each VP may have as maximum value.

For Experiment A, it can be seen that, without caching caching, VPs would only have 1 answer out of 18 sent, for the 15k+ VPs (Table 4). However, in this scenario, a median VP has 6 queries answered, largely due to effects of caching. And to retrieve these 6 answers, it takes 60 minutes (given our probing interval is 10 minutes), which is exactly the TTL of the record (3600 s, or 60 minutes).

However, for Experiments B and C, we see a different behavior: the median value of OK answers is only incremented by 1 in comparison with Max-AT, *i.e.*, if there would be no caching. This is because the DDoS starts at $t = 60m$, upon of which queries should be about to expire on most VPs (given they had TTLs of 60 m and 30 m). Thus, in this case, caching helped, but not as much. Without caching, Experiments B and C would have obtained 18 and 12 answers correctly, respectively. Thus, what we can learn from this is that operators cannot assume the TTL of their zone determines how long all their clients will be covered in case of DDoS; it actually determines the upper bound (scenario A). The state of each cache will determine that actual value.

F ANALYSIS OF RETRIES FROM A SPECIFIC PROBE

In §6 we examined changes in legitimate traffic during a DDoS event, showing traffic greatly increases because of retries by each of possibly several recursive resolvers. That analysis looked at aggregate statistics. We next examine a *specific* probe to confirm our evaluation.

We consider RIPE Atlas Probe 28477, in Experiment I. This probe is located in Australia. It has three “local recursives” ($R1_a$, $R1_b$, $R1_c$) and 8 “last layer” recursives (the ones that ultimately send queries to authoritatives, Rn_1 – Rn_8), as can be seen in Figure 16.

In this section, we analyze all the incoming queries arriving on our authoritatives for AAAA records of 28477.ca.chetest.nl. We do not analyze NS queries, since we cannot associate them to this probe.

We then analyze data collected at both the client side (Ripe Atlas probes) and the server side (Authoritative servers). We show the summary of the results in Table 6. In this table, we highlight in light color the probing intervals T in which the simulated DDoS started (dataset I, 90% packet drop on both authoritative servers – Table 4).

Client Side: Analyzing this table, we can learn the following from the Client side:

During *normal operations*, we see 3:3:3 at the client side (3 queries, 3 answers, and 3 $R1_n$ used in the answers). By default, each Ripe Atlas probe sends a DNS query to each $R1_n$.

During the *DDoS*, this slightly change: we see 2 answers from 2 $R1_s$ – the other either times out or SERVFAIL. Thus, at 90% packet loss on both authoritatives, still 2 of 3 queries are answered at the client (and there is no caching since the TTL of records is shorter than the probing interval).

This is in fact a good result given the simulated DDoS.

Authoritative side: Now, let’s consider the authoritative side:

In *normal operation* (before the DDoS), we see that the 3 queries sent by the clients (client view) lead to 3 to 6 queries to be send by the R_n , which reach the authoritatives. Every query is answered, and both authoritatives are used (# ATs).

We also see that the number of R_n used before the attack changes from 2 to 6 (# R_n). This choice depends on how $R1_*$ choose their upper layer recursives (R_{n-1}) and ultimately how R_{n-1} chooses R_n .

We also can observe how each R_n chooses each authoritative (# R_n -AT): this metric show the number of unique combinations of R_n and authoritatives observed. For example, for $T = 1$, we see that three queries reach the authoritatives, but only 2 R_n were used. By looking at the number of unique recursive-AT combination, we see that there are three. As a consequence, one of the R_n must have used both authoritatives.

During the *DDoS*, however, things change: we see that the still 3 queries send by probe 28477 (client side) now turn into 11–29 queries received at the authoritatives, from the maximum of 6 queries before the DDoS. The client remains unaware of this.

We see three factors at play here:

- (1) **Increase in number of used R_n :** we see that the number of number of used R_n also increases a little (3.6 average before to 4.5 during). (However, not all R_n are used all the time during the DDoS. We have not enough data of why this happens – but surely is how the lower layer recursives chooses these ones, or how they are set up.)
- (2) **Increase in the number of authoritatives used by each R_n :** Each R_n , in turns, contacts an average of 1.13

authoritatives before the attack (average of $UniqR_n$ -AT/ $UniqR_n$). During the DDoS, this number increases to 1.37.

- (3) **Retrials by each recursive,** either by switching authoritatives or not: we found this to be the most important factor for this probe. We show the number of queries sent by the top 2 recursives (QueriesTop2Rn) and see that compared to the total number of queries, they comprise the far majority.

For this probe, the answer of why so many queries is a combination of factors, being the retrials by R_n the most important one. From the client side, there is no strong indication that the client’s queries are leading to such an increase in the number of queries to authoritatives during a DDoS.

G ADDITIONAL INFORMATION ABOUT SOURCES OF RETRIES

In §6.2, we summarized sources of retries. Here we provide additional experiments to confirm prior work by Yu et al. [50] still holds. We next examine two popular recursive resolver implementations: BIND 9.10.3 and Unbound 1.5.8.

To evaluate these implementations we deploy our own recursive resolvers, and record traffic with the authoritatives (cachetest.nl) up and then with them inaccessible. (We ignore traffic from the recursives to other authoritatives.) For consistency, we only retrieve AAAA records from our authoritative servers. The recursives operates normally, learning about the authoritatives from [.nl](http://nl) and querying one or both based on its internal algorithms. All queries are made with a cold cache.

Figure 17 shows the results of this experiment, with normal behavior in the left two groups, and the inaccessible DDoS attempts in the right two. Each group of bars counts different types of queries that are made.

From this figure we see that, in normal operation, a few queries are required to look up the target domain: the server looks up AAAA records for both nameservers and then queries the AAAA record for the target. This request is with a cold cache, and because the AAAA-ns records do not exist. Negative information will be cached for the AAAA-ns requests, allowing those queries to be skipped for other, subsequent requests to the domain. (And the AAAA-pid query would be skipped if it was immediately re-requested.)

The two groups on the right show much more aggressive queries during server failure. Both bind and unbound make each specific request 6 or 7 times from each authoritative as it tries to get the AAAA records for each of both nameservers and the target record. This experiment is consistent with our observation that we see 7× more traffic than usual during near-complete failure.

We do not mean to suggest that these retries are incorrect—recursives need to retry to account for packet loss. Both

T	Client View (Atlas probe)			Authoritative Server View						
	Queries	Answers	# R1	Queries	Answers	# ATs	# Rn	# (Rn - AT)	Max. (% total) (Que(Rn - AT))	Queries (top-2 Rn's)
1	3	3	3	3	3	2	2	3	1 (33.3%)	2;1
2	3	3	3	5	5	2	4	5	1 (20.0%)	2;1
3	3	3	3	6	6	2	6	6	1 (16.7%)	1;1
4	3	3	3	5	5	2	2	3	2 (40.0%)	4;1
5	3	3	3	3	3	2	3	3	1 (33.3%)	1;1
6	3	3	3	6	6	2	5	5	2 (33.3%)	2;1
7	3	2	2	29	3	2	2	4	10 (34.5%)	15;14
8	3	2	2	22	6	2	7	9	7 (31.8%)	10;7
9	3	2	2	21	1	2	3	6	9 (42.9%)	16;3
10	3	2	2	11	4	2	3	4	6 (54.5%)	8;2
11	3	2	2	21	3	2	4	6	15 (71.4%)	17;2
12	3	2	2	23	1	2	8	8	15 (65.2%)	15;2
13	3	3	3	3	3	2	3	3	1 (33.3%)	1;1
14	3	3	3	4	2	2	4	4	1 (25.0%)	1;1
15	3	3	3	8	8	2	7	7	7 (87.5%)	1;1
16	3	3	3	9	9	2	3	3	6 (66.7%)	6;3
17	3	3	3	9	9	2	5	5	5 (55.6%)	5;1

Table 6: Client and Authoritative view of AAAA queries (probe 28477, measurement I)




Figure 17: Histogram of Queries from to cachetest.nl authoritative servers




Figure 18: Exponential back-off: queries timeseries from BIND and Unbound to authoritative under total failure. Points show switching from one authoritative to another.

implementations do exponential backoff in the timing of these requests – this can be seen in Figure 18. We point out that exponential backoff is important, and encourage all implementations to be careful about traffic to apparently down authoritatives.