Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event (extended)

USC/ISI Technical Report ISI-TR-2016-709, May 2016

Giovane C. M. Moura[1] Ricardo de O. Schmidt[2] John Heidemann[3]
Wouter B. de Vries[2] Moritz Müller[1] Lan Wei[3] Cristian Hesselman[1]
1: SIDN Labs 2: University of Twente 3: USC/Information Sciences Institute

ABSTRACT

Distributed Denial-of-Service (DDoS) attacks continue to be a major threat in the Internet today. DDoS attacks overwhelm target services with requests or other traffic, causing requests from legitimate users to be shut out. A common defense against DDoS is to replicate the service in multiple physical locations or sites. If all sites announce a common IP address, BGP will associate users around the Internet with a nearby site, defining the *catchment* of that site. Anycast addresses DDoS both by increasing capacity to the aggregate of many sites, and allowing each catchment to contain attack traffic leaving other sites unaffected. IP anycast is widely used for commercial CDNs and essential infrastructure such as DNS, but there is little evaluation of anycast under stress. This paper provides the *first evaluation of several anycast services under stress with public data*. Our subject is the Internet's Root Domain Name Service, made up of 13 independently designed services ("letters", 11 with IP anycast) running at more than 500 sites. Many of these services were stressed by sustained traffic at 100× normal load on Nov. 30 and Dec. 1, 2015. We use public data for most of our analysis to examine how different services respond to the these events. We see how different anycast deployments respond to stress, and identify two policies: sites may *absorb* attack traffic, containing the damage but reducing service to some users, or they may *withdraw* routes to shift both good and bad traffic to other sites. We study how these deployments policies result in different levels of service to different users. We also show evidence of *collateral damage* on other services located near the attacks.

1. INTRODUCTION

Although not new, denial-of-service (DoS) attacks are a continued and growing challenge for Internet services (for example, [2, 3]). In a DoS attack the attacker overwhelms a service, with large amounts of either bogus traffic or seemingly legitimate requests. Actual legitimate requests are lost due to limits in network or compute resources at the service. Once overwhelmed, the service is susceptible to extortion [32]. Persistent attacks may drive clients to other services. In some cases, attacks last for weeks [13].

DDoS attacks are possible because of three reasons: First, source-address spoofing allows a single machine to masquerade as many machines, making filtering difficult. Second, attackers use protocol approaches to amplify their attacks, so for each byte sent by an attacker, 5 or 20 (or more!) bytes are delivered to the victim. Third, attackers can easily get botnets of thousands of machines, so even without spoofing and amplification, vast attacks are possible. Large attacks today are in the 50–350 Gb/s range [50], and 1 Tb/s attacks are certainly within reach.

Many protocol-level defenses against DNS-based DDoS attacks have been proposed. Source-address validation prevents spoofing [20]. Response-rate limiting [46] reduces the effect of amplification. Protocol changes such as DNS cookies [17] or broader use of TCP [54] can blunt the risks of UDP. While these approaches reduce the effects of a DoS attack, they cannot eliminate it. Moreover, deployment rates of these approaches have been slow [7], at least in part because there is a mismatch of incentives between who must deploy these tools (everyone) and the victims of attacks.

Defenses in protocols and filtering are limited, though—ultimately the only defense to a 10000-node botnet making legitimate-appearing requests is capacity. Services can be replicated to many IP addresses, and each IP address can use IP anycast to operate at multiple locations. This allows single services to have a high capacity in terms of processing and bandwidth.

Many commercial services promise to defend against DDoS, either by offering DDoS-filtering as a service (as provided by Verizon, NTT, and many others), or by supporting a particular service in a DDoS-resistant way (such as Akamai, Cloudflare, and others). Yet the specific impact of DDoS on real infrastructure has not widely been reported, often because commercial infrastructure is proprietary and each service's "secret sauce".

The DNS is a common service, and the root servers are a fundamental, high-profile and publicly visible service that has been subject to DoS attacks in the past.

As a public service, they are observed [34] and strive to self-report their performance. Perhaps unique among many large services, the Root DNS service is operated by 12 different organizations, with different implementations and infrastructure. Although the internals of each implementation are not public, some details (such as the number of anycast sites) are.

To evaluate the effects of DoS attacks on real-world infrastructure, we analyze one specific event: the DNS Root events of Nov. and Dec. 2015 (see § 2.3 for discussion and references). We investigate how the DDoS attack affected reachability and performance of the anycast deployments. This paper is the first to explore the response of real infrastructure across several levels, from specific anycast services (§ 3.2), physical sites of those services (§ 3.3), and of individual servers (§ 3.5). An important consequence of high load on sites are routing changes, as users "flip" from one site to another after a site becomes overloaded (§ 3.4).

The overall contribution of this paper is the *first evaluation of several anycast services under stress with public data*. Anycast is in wide use and commercial operators have been subject to repeated attacks, some of which have been reported, but the details of those attacks are often withheld as proprietary. We demonstrate that in large anycast instances, *site failures* can occur even if the service as a whole continues to operate. Anycast can both *absorb* attack traffic inside sites, and also *withdraw* routes to shift both good and traffic to other sites. We explore these policy choices in the context of a real-world attack, and show that *site flips do not necessarily help* if the new site is also overloaded, or if the shift of traffic overloads it. Finally, we show evidence of *collateral damage* (§ 3.6) on services near the attacks. These results and policies presented can be used by anycast operators in the management of their infrastructure. Finally, the challenges we showed suggest potential future research in improving routing adaptation under stress and provisioning anycast to tolerate attacks.

2. BACKGROUND AND DATASETS

The goal of our paper is to assess anycast services under attack. We next summarize how anycast works, then describe the denial-of-service attack on Root DNS service on Nov. 30 and Dec. 1, 2015 and the datasets we use to study how it affected Root DNS anycast services.

2.1 Anycast Background and Terminology

This paper is interested in understanding how anycast services react to stress, so we next summarize how IP anycast works using the Root DNS service as an example.

Root DNS service is implemented with several mechanisms operating at different levels (Figure 1): a `root.hints` file to bootstrap, multiple IP services, often anycast; BGP routing in each anycast server; and often multiple servers at each site.

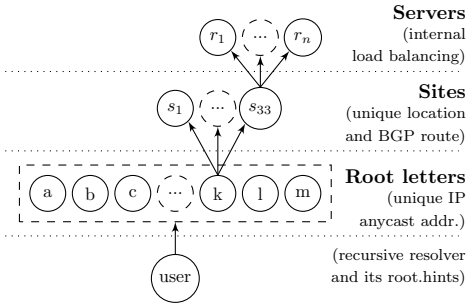

Figure 1: DNS Root structure, terminology, and mechanisms in use at each level.

letter	operator	sites (global, local)	architecture
A	Verisign	5 (5, 0)	anycast
B	USC/ISI	1 (1, 0)	single site
C	Cogent	8 (8, 0)	anycast
D	U. Maryland	87 (18, 69)	anycast
E	NASA	12 (1, 11)	anycast
F	ISC	59 (5, 54)	anycast
G	U.S. DoD	6 (6, 0)	anycast
H	ARL	2 (2, 0)	primary/backup
I	Netnod	49 (49, 0)	anycast
J	Verisign	98 (66, 32)	anycast
K	RIPE	33 (15, 18)	anycast
L	ICANN	144 (144, 0)	anycast
M	WIDE	7 (6, 1)	anycast

Table 1: The 13 Root Letters, each operating a separate DNS service, and their number of sites and architecture as of 2015-11-18 [38].

DNS Root service is implemented by 13 separate DNS services (Table 1), each running on a different IP address, but sharing a common master data source. These are called the 13 *DNS Root Letter Services* (or just the "Root Letters" for short), since each is assigned a letter from A to M and identified as `X.root-servers.net`. The letters are operated by 12 independent organizations and each letter has a different architecture, an intentional diversity designed to provide robustness. This diversity happens to provide a rich natural experiment that allows us to explore how different approaches react to the stress of common attacks.

Most Root Letters are operated using IP anycast [1]. At the time of the attack, only B-Root was unicast [38], and H-Root operated with primary-backup routing [31]. In IP anycast, the same IP address is announced from multiple *anycast sites* (s_1 to s_{33} in Figure 1), each at a different physical location. BGP routing associates clients (users) who chose to use that service with a nearby anycast site. The set of users of each site defines its *anycast catchment*.

Larger anycast sites may consist of multiple physical servers (r_1 to r_n in Figure 1), each an individual machine that responds to queries. Sites and servers can often have unique responses to CHAOS queries [51], but replying to these queries is optional (and queries can be spoofed by third parties). While the format of replies are not standardized, most letters follow patterns that they disclose or can be inferred. (Prior studies have used traceroute and other approaches to detect masquerading [19].)

Root Letters have different policies, architectures, and sizes, as shown in Table 1. Some letters constrain routing to some sites to be *local*, using BGP policies (such as `NOPEER` and `NO_EXPORT`) to limit routing to that site to only its immediate or neighboring ASes. Routing for *global* sites, by contrast, is not constrained.

2.2 Anycast vs. DDoS: Design Options

How should an anycast service react to the stress of a DDoS attack? A site under stress, overloaded with incoming traffic, has two options. It can *withdraw* routes to some or all of its neighbors, shrinking its catchment and shifting both legitimate and attack traffic to other anycast sites. Possibly those sites will have greater capacity and service the queries. Alternatively, it can be become a *degraded absorber*, continuing to operate, but with overloaded ingress routers, dropping incoming legitimate requests due to queue overflow. However, continued operation will also absorb traffic from attackers in its catchment, protecting other anycast sites [1].

These options represent different results in anycast deployments. A withdrawal strategy causes anycast to respond as a waterbed mattress, with queries displaced from one site shifting to others. The absorption strategy is a conventional mattress, "compressing" under load, with queries getting delayed or dropped. We see both of these behaviors in practice and observe them through site reachability and RTTs next.

Although we describe these as strategies and policies, it is important to note that they are actually the *result* of the combination of operator and host ISP routing policy, routing implementations withdrawing under load [44], the nature of the attack, and the locations of the sites and attackers. Some policies are explicit, such as the choice of local-only anycast sites, or operators removing a site from service for maintenance. However, under stress, the choices of withdrawal and or absorption more often are *emergent* results of a mix of explicit choices and implementation details, such as BGP timeout values. We speculate that more careful, explicit management of policies may provide stronger defenses to overload, an area of future work.

Policies in Action: We can illustrate these policies with the following thought experiment. Consider the anycast system in Figure 2, it has three anycast sites:

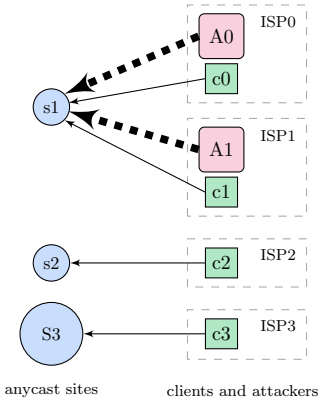

Figure 2: An example anycast deployment under stress.

s_1, s_2, S_3, four clients c_0 and c_1 in s_1's catchment, with c_2 in s_2 and c_3 in S_3's. Let A_0 represent both the identity of the attacker and the volume of its attack traffic, and s_1 represent the site and its capacity.

The best choice of defense depends on the relative sizes of traffic in A_0 and A_1 compared to the capacity of s1 and other sites. Assume $s_1 = s_2$ and $S_3 = 10s_1$. We compare alternatives mesauring how many clients are successful (H, "happiness").

1. If $A_0 + A_1 < s_1$, then the attack does not hurt anyone, $H = 4$.

2. If $A_0 + A_1 > s_1$ and $A_0 < s_1$ (and $A_1 < s_2$), then s_1 is overwhelmed ($H = 2$) but can shed load. If it withdraws its route to ISP1, A_1 and c_1 shift to s_2 and all clients are served: $H = 4$.

3. If $A_0 > s_1$ and $A_0 + A_1 < S_3$, then a attackers can overwhelm a small site, but not the bigger site. *Both* s_1 and s_2 should withdraw all routes and let the large site S_3 handle all traffic, for $H = 4$.

4. If $A_0 + A_1 > S_3$, the attack can overwhelm *any* site; the optimal strategy is to make no changes. s_1 becomes a degraded absorber and protects the other sites from the attack, at the cost of clients c_0 and c_1. $H = 2$.

This thought experiment shows that for small attacks, the *withdraw* policy can improve service by spreading the attack (less can be more!). For large attacks, *degraded absorbers* are necessary to protect some clients, at the cost of others. In practice, one cannot directly apply these rules: attack traffic volumes are unknown, because they exceed capacity; attack locations are unknown, due to source address spoofing; the effects of route changes difficult to predict, due to unknown attack locations; and route changes difficult to implement, since routing involves multiple parties. In fact, the absorption policy is likely the best choice in the face of

inevitable uncertainty about an attack's true size and location. However, route withdrawals may occur due to BGP session failure, so both policies may occur.

Another strategy employed by many commercial services is to use commercial traffic scrubbing services. Such services capture traffic using BGP then filter attack traffic and forward the clean traffic to its final destination. While cloud-based scrubbing services have been used by web companies (for example, in the ProtonMail DoS [33]), to our knowledge Root DNS providers do not use such services, likely because Root DNS traffic is a very atypical workload.

2.3 The Events of Nov. 30 and Dec. 1

On November 30, from 6:50 to 9:30 (UTC), then again on December 1, 2015 from 5:10 to 6:10, many of the of the Root DNS Letters experienced an unusual high rate of requests [39]. Traffic rates peaked at about 5M queries/s, at least at A-Root [47], more than 100× normal load. We sometimes characterize these event as an "attack" here, since a sustained traffic of this volume seems unlikely to be accidental, but the target of these events is unclear.

An early report by the Root Operators stated that several letters received high rates of queries for 160 minutes on Nov. 30 and 60 minutes on Dec. 1 [39]. Queries used fixed names, but source address were randomized. Some letters saw up to 5 million DNS queries per second, and some sites at some letters were overwhelmed by this traffic, although several letters were continuously reachable during the attack (either because they had sufficient capacity or were not attacked). There were no known reports of end-user visible errors, because top-level names are extensively cached, and the DNS system is designed to retry and operate in the face of partial failure.

A subsequent report by Verisign, operator of A- and J-Root, provides additional details [47]. They stated that it was limited to IPv4 and UDP packets, and that D-, L-, and M-root were not attacked. They confirm that the event queries used fixed names, with `www.336901.com` on Nov. 30 and `www.916yy.com` on Dec. 1. They reported that A and J together saw 895M different IP addresses, strongly suggesting source address spoofing, although the top 200 source addresses accounted for 68% of the queries. They reported that both A- and J-Root were attacked, with A continuing to serve all regular queries throughout, and J suffering a small amount of packet loss. They reported that Response Rate Limiting was effective, identifying duplicated queries to drop 60% of the responses, and filtering on the fixed names was also able to reduce outgoing traffic. They suggested the traffic was caused by a botnet.

Motivation: We do not have firm conclusions about the motivation for these events. As Wessels first observed [49], the intent is unclear. The events do not appear to be DNS amplification to affect others since the spoofed sources spread reply traffic widely. They might be a DDoS targeted at the services at the fixed names listed above, but `.com` must resolve those names, not the roots. Even if so, caching the results (and not bothering the roots) would provide a better attack on those targets if that were the intent. Possibly it was an attack on those targets that went awry due to bugs in the attack code. It may be a direct attack on the DNS root, or even a diversion from another attack.

Fortunately, the intent of the event is irrelevant to our use of the event to understand anycast systems under stress.

2.4 Datasets

We use these large events to assess anycast operation under stress. Our evaluation uses publicly available datasets provided by RIPE, several of the Root Operators, and the BGPmon project. We thank these organizations for making this data available to us and other researchers. We next describe these data sources and how we analyze it. The resulting dataset from the processing described next is publicly available at `http://traces.simpleweb.org/`.

2.4.1 *RIPE Atlas Datasets*

RIPE Atlas provides a measurement platform with more than 9000 globally distributed Atlas Probes that provide *vantage points* (or just *VPs*) that conduct network measurements [24, 36, 37]. For our study, the essential aspect of RIPE Atlas is that all Atlas nodes regularly probe all DNS Root Letters. A portion of this data appears in RIPE's DNSMON dashboard of Root DNS response [34]. We employ the same set of VPs for each root letter, having a distinct measurement ID per root letter [35]. We use the more detailed raw data they make public [35], with new processing as we describe below.

RIPE's baseline measurements send a DNS CHAOS query to each root letter every 4 minutes. At the time of the event, A-Root was an exception and ws probed only every 30 minutes, too infrequent for our analysis (§ 3.2) (it is now probed as frequent as the other letters.) Responses to CHAOS queries are specific to root letters (after cleaning, described below) but each letter follows a pattern that can be parsed to determine the site and server that VP sees. For this report we normalize identification of roots in the format *X-APT*, where *X* is the Root Letter (A to M) and *APT* is a three-letter airport code near the site.

Data cleaning: We take several steps to clean RIPE data for using it in our analysis. Cleaning preserves nearly all VPs (more than 9000 of the 9363 currently active in May 2016), but discards data that otherwise

would complicate analysis or provide outliers. We keep only data from VPs with firmware version 4570 or above to exclude older VPs, in order to have a more similar set of probes/firmware version.

We discard measurements from a few VPs (67, less than 1%) where traffic to a root appears to be hijacked and served by third parties. We identify hijacking manually by unusual CHAOS replies, confirmed with unusually short RTTs (less than 7 ms), following prior work [19].

After cleaning we map all observations into a time series with ten-minute bins. In each time bin we identify, for each root letter, the response: either a site the VP sees, an response error code [30], or a absence of a reply. These analysis bins each represent 2.5 RIPE probing intervals, allowing us to synchronize RIPE measurements that otherwise occur at arbitrary phases. (When we have differing replies in one bin, we prefer sites over errors, and errors over missing replies.)

Limitations of RIPE Atlas: RIPE Atlas has known limitations: although VPs are global, their locations are heavily biased towards Europe. Because they measure specific DNS letters, they do not represent "user" queries (although this difference is necessary for our analysis). VPs fail independently. We account for uneven VP location by considering data for anycast sites with sufficient observers, requiring a median of 20 VPs per catchment over the two days we study.

2.4.2 *RSSAC-002*

RSSAC-002 is a specification for data collection about Root DNS service [41]. It provides daily query rates, distributions of query sizes, and other operationally-relevant data.

All Root services have committed to provide RSSAC-002 data by the end of 2016. At the time of the events, only five services (A, H, J, K, and L) were providing this data (see RSSAC links at `http://www.root-servers.org`). In addition, RSSAC-002 monitoring is a "best effort" activity that is not considered as essential as operational service, so reporting may be incomplete, particularly at times of stress.

2.4.3 *BGPMon*

We use BGP routing data from BGPmon [52]. BGPmon has peers to dozens of routers providing full routing tables from different locations around the Internet. We use data from 152 of them to evaluate route changes at anycast sites in § 3.4.1.

3. ANALYSIS OF THE EVENTS

We next look at our analysis of the events. We begin with overall estimates on event sizes, then drill down on specific Root Letters, then anycast sites for some letters, and the individual servers at those sites. We then reconsider the attack as a whole and its effects on other services.

3.1 How big was the event?

We next estimate the size of the events, using RSSAC-002 reports, and assuming all affected letters were affected equally.

We first look at RSSAC-002 data. As we previously described (§ 2.4.2), RSSAC-002 data is best-effort, and we expect it to be incomplete.

RSSAC-002 statistics are reported per day, so to estimate the event size we took seven days before event as normal traffic, then looked at what changed on the two event days. (We omit one outlier for A-Root to avoid bias due to an independent attack on 2015-11-28 and scale appropriately.) Query sizes are reported in bins of 16 bytes. We approximate attack query and response sizes by looking for bins that were outliers on the event days. We estimate that queries were between 32 and 47 bytes on Nov. 30 and 16 and 31 bytes on Dec. 1; response sizes were between 480 and 495 bytes for both events. These sizes are for DNS payload only; to them we need to add headers. Finally, Verisign stated that the attacks were of specific query names (see § 2.3), so we generated queries with these names to confirm sizes with headers of 84 and 85 bytes for queries and 493 or 494 bytes for responses, consistent with RSSAC-002 reports. We use these approximate sizes to estimate attack bitrate.

Table 2 gives our estimates on event traffic from the five letters reporting RSSAC-002 statistics. These reported values differ greatly across letters and between queries and responses. We believe differences across letter represent measurement error, with most letters under-measuring traffic when under attack. (Under-reporting is consistent with large amounts of lost queries described in § 3.2.) The lower number of responses may be due to Response Rate Limiting [46], suppressing duplicate queries from the same source address [49]. In addition, our lower estimate ignores letters not providing RSSAC-002 data.

We propose an upper-bound for event size by correcting for both of these types of under-report. First, we assume that A-Root's RSSAC-002 data measured the entire event. Verisign reported A-Root graphs of input traffic showing about 5Mq/s at both A- and J-Root (although J's RSSAC-002 reports are much lower). Second, Verisign reported that 10 of 13 letters were attacked (D, L and M were not attacked). Finally, we assume that all attacked letters received equal traffic. We believe the first two assumptions are well justified. The third is speculation, however it seems unlikely that a well-provisioned A-Root would receive less traffic than other letters. Our upper-bound for event size is therefore ten-times A-Root's traffic.

RSSAC reports	2015-11-30 (160 min.) queries Mq/s	Gb/s	M IPs (ratio)	responses Mq/s	Gb/s	2015-12-01 (60 min.) queries Mq/s	Gb/s	M IPs (ratio)	responses Mq/s	Gb/s	Baseline queries Mq/s	M IPs
A	5.12	3.44	1,813.38(339.9x)	3.84	15.13	5.21	3.54	1345.46(252.4.0x)	3.93	15.53	0.04	5.35
H	0.23	0.15	36.14(13.2x)	0.00	0.00	0.32	0.22	16.22(6.5x)	0.00	0.00	0.03	2.94
J	1.90	1.28	18,268.5(276.6x)	1.10	4.32	2.29	1.56	8,236.6(128.1x)	1.43	5.66	0.05	2.78
K	1.07	0.72	39.23(14.4x)	0.48	0.32	1.12	0.76	40.88(15.0x)	0.28	1.09	0.04	2.92
L	0.05	0.04	36.15(13.2x)	0.05	0.19	0.10	0.07	16.22(6.5x)	0.09	0.37	0.06	2.94
bounds (lower and *upper*):												
low.	8.38	5.63	–	5.45	19.95	9.03	6.14	–	5.71	23.12	0.22	–
up.	*51.22*	*34.42*	–	*38.37*	*151.31*	*52.09*	*35.42*	–	*39.31*	*155.35*		

Table 2: RSSAC-002 reported estimates of event traffic for all letters with RSSAC-002 data (IPv4/UDP). The lower-bound estimate (*low*) sums this reported traffic, with an upper-bound estimate (*up* assumes 10 letters were sent as much traffic as A-Root observed).

If our upper-bound estimate is correct, the aggregate size of this attack across all letters is about 35–40 Gb/s. Although attacks exceeding 100 Gb/s have been demonstrated since 2012 [2, 50], such large attacks are usually after amplification [40], as is seen in the reply traffic with an upper bound of 180 Gb/s. Directly sourced traffic of 35 Gb/s on the roots therefore represents a large attack.

We can also see for all letters a large increase (by a factor 6.5× to 339.9×) in the number of unique IP addresses (IPv4) observed by each letter during the attacks. This observation conforms with the initial reports on using of IP spoofing during these attacks [49].

3.2 How were individual Root DNS letters affected?

We next consider how each letter reacted to the event, based on observations from RIPE Atlas. We then discuss overall DNS Root performance—performance of specific letters does not reflect end-user performance, since users employ caching and will employ multiple letters if they need to refresh their cache.

3.2.1 *Reachability of specific letters*

Figure 3 shows the reachability for each root letter as measured by RIPE Atlas. We plot D-, L-, and M-Root together because overall they see no visible change, consistent with reports that they were not attacked [47]. However, we show in § 3.6 that while D-Root, as a whole, sees no change, a small number of D-root sites see collateral damage from the event. On these days Atlas probed A-Root less frequently than other letters (§ 2.4.1), so this graph we scale A's observations to account for this difference. Because infrequent probing of A-Root makes the event dynamics impossible to discern, we omit A-Root from analysis in the rest of the paper.

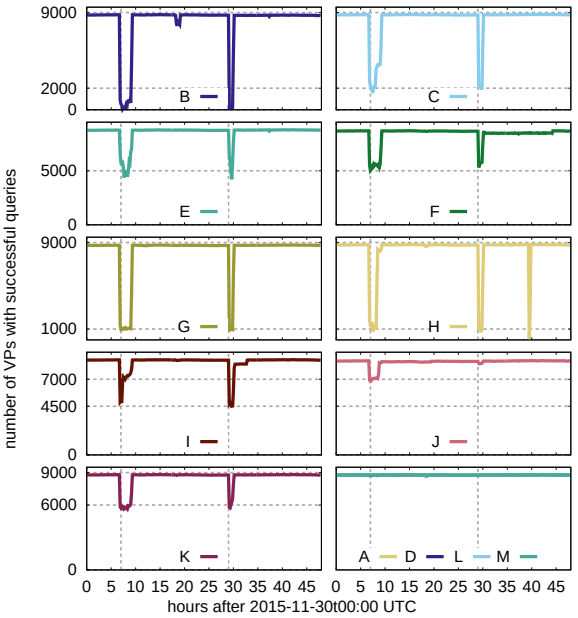

Figure 3: Number of VPs with successful queries (RCODE 0, measured in 10-minute bins). (The scales on both axes are consistent across all plots, with nearly 9000 VPs across 48 hours of observation. In all graphs, dotted lines highlight approximate event start times. Here they also show lowest values for the dips.)

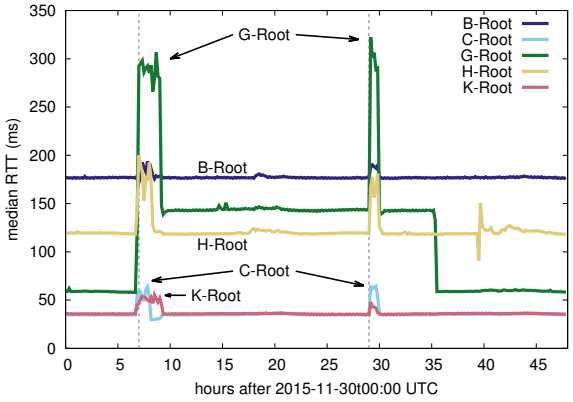

Figure 4: Median RTT for some letters during the attacks. Letters with no significant change (A, D, E, F, I, J, L and M) are omitted.

All the other letters experience different degrees of reachability problems at most sites, during the exact reported attack intervals (§ 2.3). We can see a rough correlation ($R^2 = 0.75$) between the number of sites associated with a letter and its worst responsiveness, measured by the smallest number of Atlas VPs receiving service during the events. B-Root, a unicast letter, suffered the most, followed by H, with two sites and primary-secondary routing. With many sites, J-Root sees some VPs loose service, but only a few. We evaluate the *causes* for service loss in § 3.3.

We can also evaluate overall performance for each letter by the RTT of successful queries, as shown in Figure 4. Note that each letter has a different baseline RTT, corresponding with the median distance from Atlas VPs to anycast sites for that letter. B-Root shows little change in RTT, while G- and H-Root see large changes in latency. In the next section we show that anycast sites can *fail*, causing routing to shift their traffic to other locations. Thus we believe these shifts in RTT indicate route withdrawals for sites that shift VP traffic to more distant sites. For example, H-Root has sites on the U.S. East and West coasts (north of Baltimore, Maryland, and in San Diego, California). Most Atlas VPs are in Europe, so we infer that the primary site for H is the U.S. East coast, but when that route is withdrawn (during both events) traffic shifts to the west coast. This assumption is confirmed by H's median RTT at that time matching B-Root's RTT, since B-Root is also on the U.S. West coast. We examine site route withdrawals in more detail in § 3.3.

3.2.2 *Reachability of the DNS Root as a whole*

While we see individual letters show degraded responsiveness under stress, the DNS protocol has several levels of redundancy, and a non-response from one letter should be met by a retry at another letter. This paper does not evaluate overall responsiveness of the DNS root, but our per-letter analysis shows some evidence of this redundancy.

L-Root was not subject to this attack, yet Table 2 shows that L-Root shows a significant increase in query rate during the second event, with a two-thirds increase in queries-per-second. More impressive, it sees a 6- or 13-fold increase in number of unique IPs on both event dates. We later describe "site flips", where VPs change anycast sites (§ 3.4.1); this coarse data suggests *letter flips* also occur (typically resolvers switching from one letter to another [53] due to shorter RTTs). While not the focus of this paper, these letter flips show the multiple levels of resilience in the Root DNS system.

3.3 How were anycast sites affected?

From the overall responsiveness of individual letters (§ 3.2), we now look for reasons why different letters show different response. Anycast services are composed of multiple sites (Table 1), so we next look at how the event affected specific sites inside each letter. As discussed in § 2.2, anycast operators and their hosting ISPs can design sites to withdraw routes or continue as degraded absorbers when under stress. We next look for evidence of these policies in the event.

3.3.1 *Site Reachability*

We first consider site reachability: how many VPs reach a letter's sites over the 2-day observation, measured in each ten-minute bin. The median number of VPs over the observation provides a baseline of "regular" behavior, calibrating how RIPE Atlas maps to a given service. Atlas coverage is incomplete, with sites have zero or a few VPs to thousands of VPs in their typical catchment. Our use of median normalizes coverage to identify trends, such as if the site adds or loose VPs. Addition of VPs to a site indicates that some other sites under stress withdraw. Reduction in VPs indicate that either that site withdrew some or all routes, or that it was overloaded and simply lost queries—reduction can therefore be caused by both withdrawal and absorption.

Figure 5 evaluates all sites for two letters (E- and K-Root). The numbers in between parenthesis show the median VPs at each site, while the blue lines show how much that site shrank or grew over the two days, normalized to the median.

We see that sites show two responses indicating reduced capacity. Some (such as E-AMS, K-LHR) become completely unavailable, as shown by the minimum dropping to zero. K-Root confirmed unavailability of some sites [55]. Others (E-NRT, K-WAW) become partially available, with the minimum dropping, but not to zero.

In addition, several sites show an increase above median over the period (the maximum blue value is greater than 1). Most of the well-observed K-Root sites show

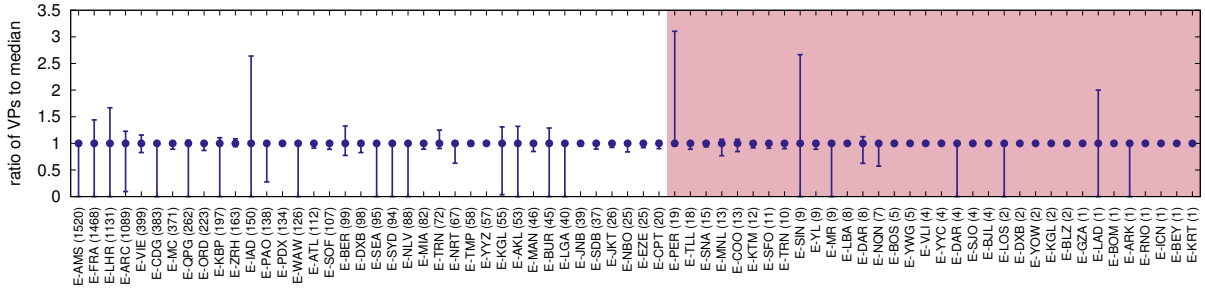

(a) E-Root sites

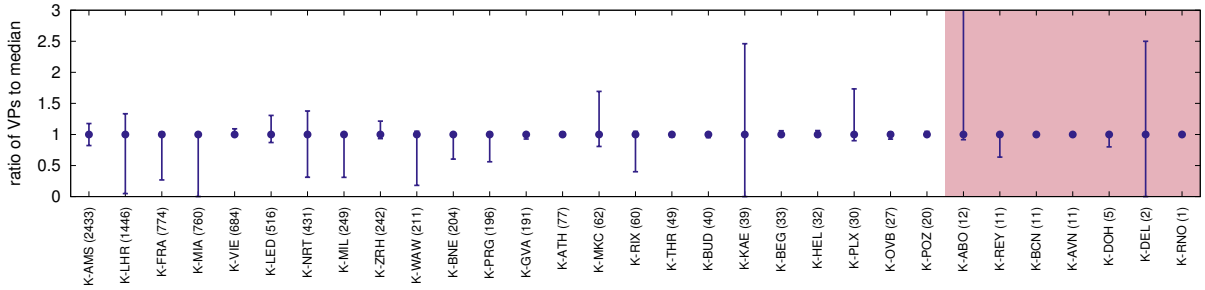

(b) K-Root sites

Figure 5: Minimum and maximum number of VPs, normalized to median (shown between parenthesis per individual site), for sites from E- and K-Root. Data for 10-minute bins over two-days. Sites are ordered by median VPs and red area highlights sites with less than 20 VPs (rule-of-thumb threshold).

some increase (K-AMS, K-LHR, K-LED, K-NRT), as do many of the well-observe E-Root sites (E-FRA, E-LHR, E-ARC, E-VIE, E-IAD).

To explore how sites behave over time, Figure 6 shows sites for E- and K-Root as time series. Each mini-plot is one site, with the line showing how many VPs are mapped to it relative to the site's median. From this figure we see that sites from these two letters behaved completely different. While most sites of E-Root either see an increase or a decrease on their reachability, most sites of K-Root seem to overlook the attack. (Note that large increases observed for few sites, such as E-DXB and K-DEL, are caused by a very low median (two VPs)—any additional VP hitting this sites during the attack can cause a peak on reachability.)

Figure 6 shows that 5 sites from E Root (E-AMS, E-CDG, E-WAW, E-SYD and E-NLV) seem to "shut down" after the attack of Dec 1 (hour 29). These sites also had reachability strongly compromised during the first event on Nov. 30 (hour 7).

What is interesting to see for the sites of both letters in Figure 6 is that major sites had reachability compromised; that is, sites with higher medians. An exception is K-AMS, with a high median, and that even accommodated more traffic than usual during the whole period.

The increase over median observed mostly for sites of E-Root suggest we are seeing some examples of route withdrawals at other sites. However, that does not explain why letters show reduced overall reachability (Figure 3), since if overloaded sites fail and traffic shifts all queries should be answered. We next look for evidence of degraded absorption.

3.3.2 *Site RTT Performance*

Sites that remain accessible may also be overloaded. RTT of successful queries provides a way to assess such load.

Figure 7 shows the median RTT for selected K-Root sites. Although the K-AMS site remained up and showed minimal loss, its median RTT showed a huge increase: from roughly 30 ms to 1 s on Nov. 30, and to almost 2 s on Dec. 1, strongly suggesting the site was overloaded. K-NRT shows similar behavior, with median RTT rising from 80 ms to 1 s and 1.7 s in the two events. Overload does not always result in large latencies. B-Root (a single site) showed only modest RTT increases (Figure 4), since only few probes could reach it during the attack (Figure 3). We hypothesize that large RTT increases are the result of an overloaded link combined with large buffering at routers (industrial-scale bufferbloat [23]).

3.4 How can services partially fail?

We know that letters report different amounts of service degradation (Figure 3), and that their sites seem to follow two policies under stress. We next look at service reachability from a client perspective to understand how

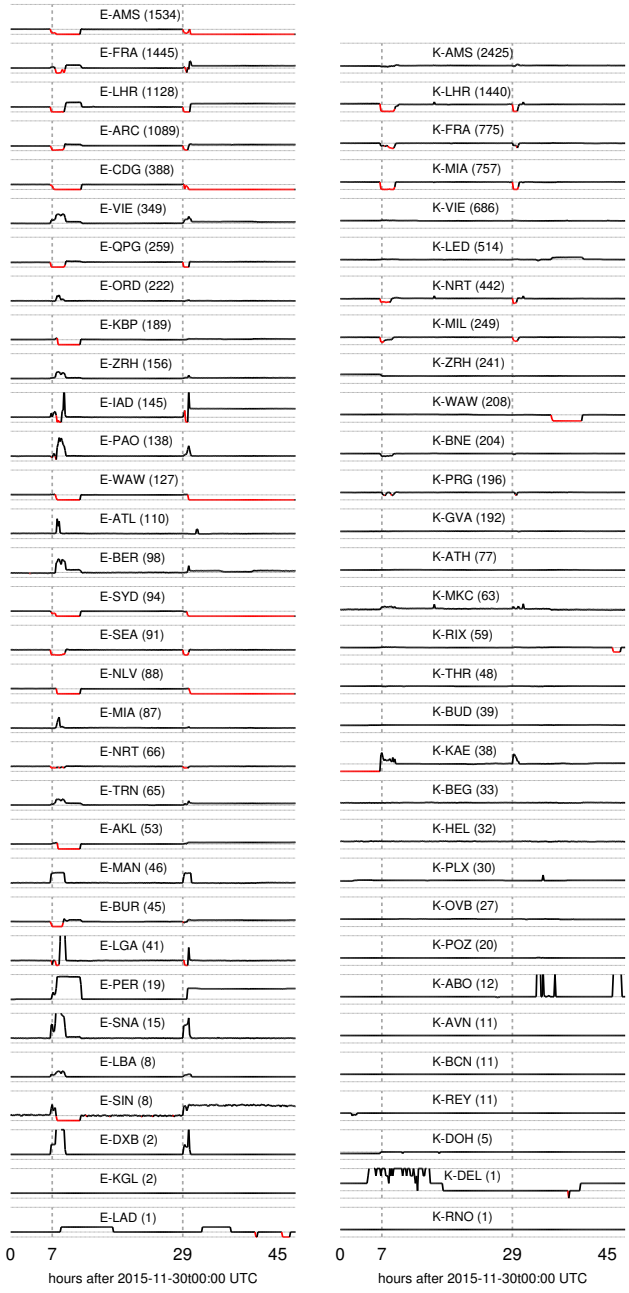

(a) E-Root sites (b) K-Root sites

Figure 6: Reachability seen by VPs that received positive responses (RCODE 0) for sites of two letters. The central line in each plot is the median, with the lower line 0 and the upper line 5× and 3× the median for E- and K-Root respectively. The red lines indicate potential *critical* moments in which reachability was dropping below the median. Sites are sorted by median (which is given between parenthesis).

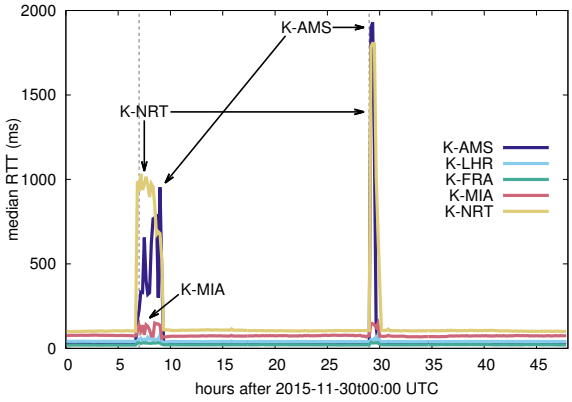

Figure 7: Performance for selected K-Root sites.

services can partially fail and what implication that has for users.

3.4.1 *Site Flips: Evidence of Stress*

A design goal of DNS and IP anycast is that service is provided by multiple IP addresses (DNS) and sites (anycast). Through their recursive resolvers, clients can turn to service on other IP address (other DNS root letters), and through BGP, to other anycast sites. A recursive DNS resolver will automatically retry with another nameserver if the first does not respond, which is intentional redundancy in the protocol and operational best practices [18, 53]. Redundancy *inside* most letters depends on IP anycast, and the routing policies DNS service operators establish at each anycast site (withdraw or absorbing).

To study a client's view of IP anycast redundancy we look for changes in site catchments. We measure these as *site flips*: when a VPs changes from its current anycast site to another. We expect each VP to have a preferred site (hopefully with low RTT), and site flips to be rare, due to routing changes or site maintenance.

Figure 8 shows site flips measured in RIPE Atlas VPs, with bursts of site flips during the event periods for letters that saw event traffic. All letters see thousands of site flips during the event (note the scale of y-axis), with E, H and K seeing many flips while C, I and J less. We next consider K-Root as a case study to show what site flips mean in practice.

To evaluate if these site flips are actually due to route withdrawals we use route data from BGPmon (§ 2.4.3). These BGPmon VPs are in different locations from our RIPE Atlas VPs, so we do not expect them to see exactly the same results, but we do expect to see more routing activity during the events.

Figure 9 shows the route changes we see. With BGPmon VPs and root anycast sites around the world we see occasional route changes over the whole time period. With 152 VPs, a routing change near one site can

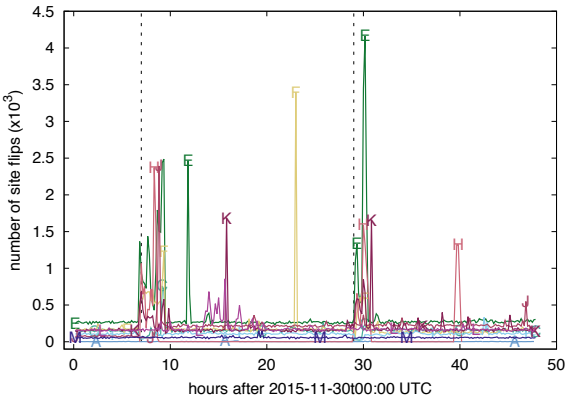

Figure 8: Number of site flips per root letter.

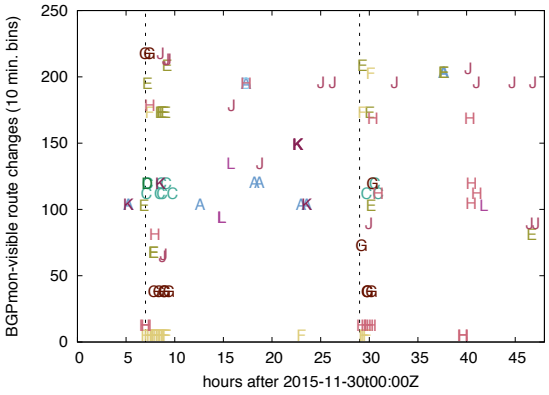

Figure 9: Route changes for each root letter (10 minute bins, seen from BGPmon route collectors).

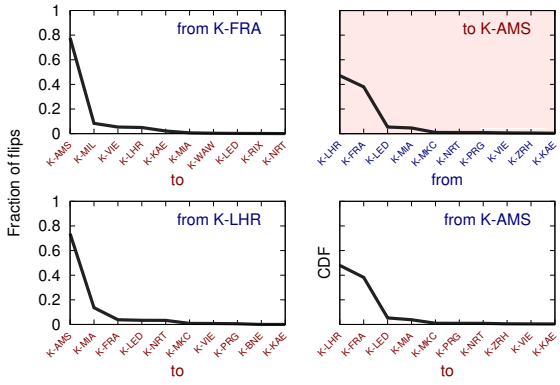

Figure 10: Site flips for selected K-Root instances over the two days.

often be seen at 100 or more VPs. But the *very frequent* sets of changes shown by *many* letters in the two event periods (4 to 6 hours and around 29 hours) suggests event-driven route changes for many letters (C, E, F, G, H, J, K). Route changes for K-Root do not appear at our BGP observers for the second event, and K's BGP changes are lower than we expect based on site flips. We suspect that our BGP vantage points are U.S.-based, while site flips are VPs that are much more numerous in Europe.

3.4.2 *Case study: K-Root*

K-Root's sites provide good examples of different policies under stress. We next consider VPs that start at K-LHR and K-FRA (London and Frankfurt), two European sites that lost nearly all or about half of VPs during the event, to see what happened to these clients. From Figure 3 we know that some clients were unsuccessful, while the maximums in Figure 5b show that some sites gained clients.

Figure 10 shows where sites from K-LHR and K-FRA went over the measurement period—the left two graphs show that about 80% of all VPs that shifted traffic during the events shift to K-AMS (Amsterdam). The top right gray graph shows where new VPs that see K-AMS just were, confirming they arrive from K-LHR and K-FRA. The bottom right gray shows that K-AMS sites also shift back to K-LHR and K-FRA as their preferred catchments after the events.

However, a question persists: *if traffic shifts to other sites and K has excess capacity, why do some VPs fail to reach K during the attack?* The reason is the dynamics that result from routing policies and implementation details (§ 2.2) at each site and its hosts: those policies and details can result in a site that will continue to receive traffic from its peer and operate as a degraded absorber, or that will withdraw its route and reallocate

its catchment. We see evidence of both results in K-Root.

To demonstrate these policies at work we must look at the actions of individual VPs. Figure 11 shows 300 randomly selected VPs that start at K-LHR (yellow) and K-FRA (salmon) for 36 hours. Each pixel represents the site choice of that VP in 4-minute bins. Black indicates the VP got no reply, while blue and white indicate selection of K-AMS or some other K-Root site.

We focus on the 40 VPs shown in Figure 11b and see two behaviors during the event and three after. During the event the top 10 VPs (labeled (1)) stick to K-LHR, but only get occasional replies. They represent a degraded absorbing peering relationship; these clients seem "stuck" to the K-LHR site. The next group labeled (2) shift to K-FRA (salmon) during the event and for a short period after, then return to K-LHR. However, during their visit to K-AMS only about a third of their queries are successful. This group shows that K-AMS is overloaded but up, and that these VPs are in ASes that are not bound to K-LHR. For the third group, marked (3), some stay at K-LHR during the event, while others shift to other sites, but all find other sites after the event. Finally, the group (4) shifts to K-FRA during the event and remains there afterward. We see similar groups for the K-FRA sites in the first event and for both sites in the second event.

We believe this kind of *partial failure* represent a *success* of anycast in isolating some traffic to keep other sites functional, but this degraded absorbing policy results in some users suffering during the event due to the overload at K-LHR. While this policy successfully protects most K-Root sites during the event, it also suggests opportunities for alternate policies during attack. Rather than let sites fail or succeed, services may choose to control routing to engineer traffic to provide good service to more users. Alternatively, if attack traffic is localized, services may choose to target routing so that only one catchment is affected—a policy particularly appropriate for attacks where all traffic originates from a single location, even if it spoofed source addresses.

3.5 How were individual servers affected?

Large anycast sites may be provisioned by multiple servers behind a load balancer (Figure 1). We now focus on the effects of the events on individual servers within specific anycast sites. We look at two sites of K-Root, K-FRA and K-NRT as examples. Although we saw similar behavior at some other sites, we do not categorize how all sites or servers behave.

Figure 12 examines which servers at K-FRA (top) and K-NRT (bottom) respond when faced with high demand during the events. At K-FRA, we typically saw replies from three servers. As the load of each event rose replies shifted to come from only one server, with none from the other two we previously saw replying. Which server responded was different in the two events, with K-FRA-S2 replying in the first event and -S3 in the second. We do not know if the other two servers failed, or if they were only serving attack traffic, or if traffic from these VPs was somehow isolated from attack traffic. Either way, this strategy seems to work reasonably well since Figure 13 shows that, after a short increase in RTT at the beginning of the attack, the median RTT for K-FRA remains stable for successful replies throughout the attack. However, K-FRA seems to be overloaded and dropping queries, as shown in Figure 6b and Figure 11b.

K-Root's Tokyo site (K-NRT) shows a different result. Figure 12 (bottom) shows that VPs have difficulties to reach all the three servers from K-NRT during the event. This suggest that the event was affecting all K-NRT servers, either because load balancing was mixing our observations with attack traffic, or because attack traffic was congesting a shared link. Figure 13 (bottom) shows larger latencies for successful queries at K-NRT, perhaps suggesting queueing at the router. We also observe that K-NRT-S2 seems more heavily loaded than the other two servers at K-NRT.

These two examples show the influence of middleboxes (load balancers) on anycast operation, in addition to BGP routing.

3.6 Did the attacks cause collateral damage?

Many Root DNS servers are placed in data centers that are shared with other services. These services may be unrelated, other infrastructure (such as other top-level domains, TLDs), even other Root DNS sites. Co-locating services creates some degree of shared risk, in that stress on one service may spill over into another causing *collateral damage*. Although data centers and operators strive to minimize collateral damage through redundancy and overcapacity, prior examples of show that it occurs and is one factor in DDoS extortion [33].

Hosting details are usually considered proprietary, and commonality can exist at many layers, from the physical facility to peering to upstream providers, making it difficult to assess shared risk.

Here we assess shared risk by end-to-end evaluation: we look for service problems in other services not directly the target of event traffic. We study two services: D-Root, a letter that was not directly attacked [47], and the `.nl` TLD. Although collateral damage is a common side-effect of DDoS, prior reports describe it as a problem but provide only few details [33].

D-Root: Figure 14 shows the absolute counts of number of RIPE Atlas VPs that reach several D-Root sites. D-Root has many sites; we report only subsets that had at least a 10% decrease in reachability during

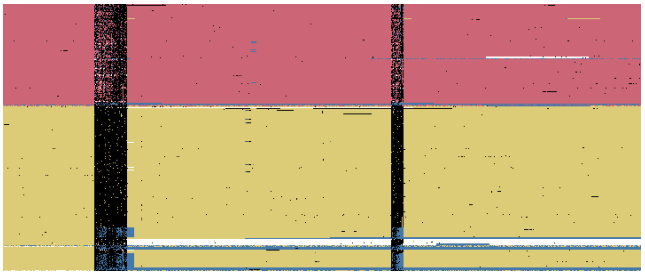

(a) A sample of 300 VPs; start 2015-11-30t00:00Z for 36 hours.

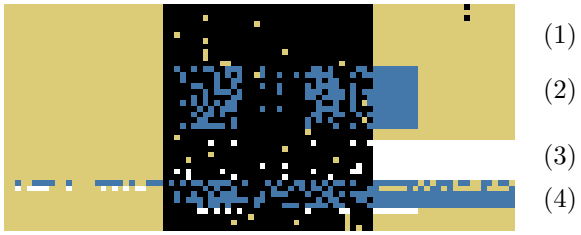

(b) A samller sample: 40 K-LHR-preferring VPs around the first event.

Figure 11: A sample of 300 VPs for K-Root that start at K-LHR (yellow) and K-FRA (salmon), with locations before, during, and after attacks. Other sites are K-AMS (blue), with white indicating other K sites, and black fail on getting a response (timeout). Dataset: RIPE Atlas.

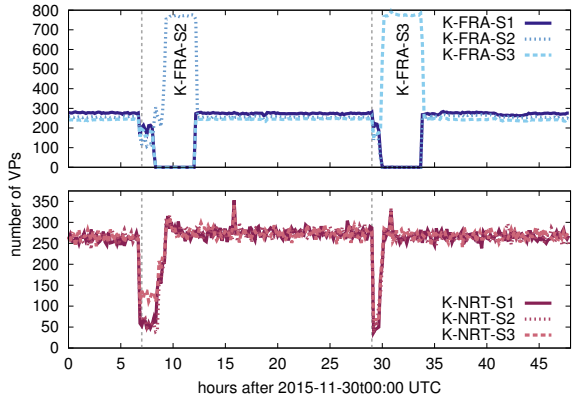

Figure 12: Reachability for individual servers from K-FRA (top) and K-NRT (bottom).

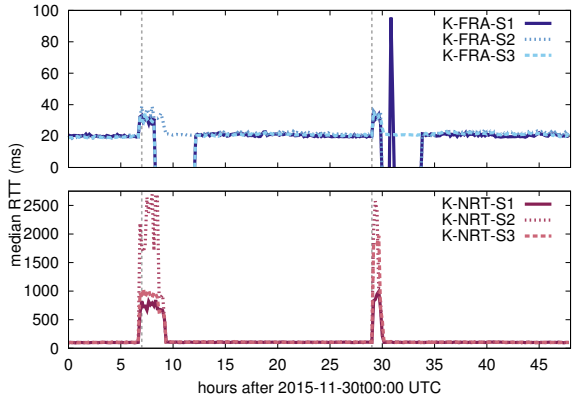

Figure 13: Performance for individual servers from K-FRA (top) and K-NRT (bottom).

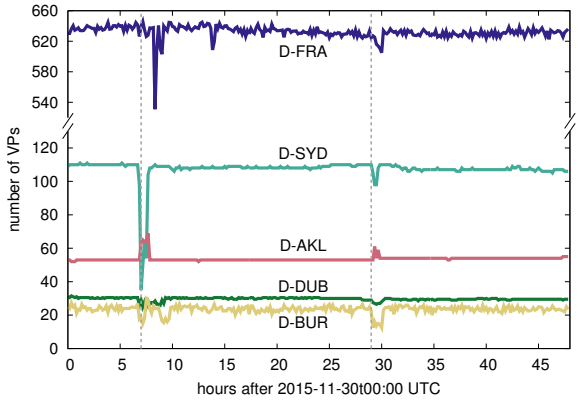

Figure 14: Reachability of those D-Root sites that were affected by the DDoS.

the time of the attacks and were reached by at least 20 RIPE Atlas probes.

These figures show the D-FRA and D-SYD sites both lost VPs during the event. The exact hosting locations of these sites is not public, but the correlation of these changes with the events suggests potential collateral damage.

Frankfurt: There are seven root letters hosted in Frankfurt (A, C, D, E, F, I, and K), and we previously observed that traffic shifted to K-FRA and yet that site suffered loss (§ 3.4.2).

D-FRA sees only small shifts in traffic, suggesting it was only slightly affected by the attacks on other letters' sites in the same city. However, this change suggests some collateral damage to D-FRA from the event.

The .nl Top-Level Domain: Finally, we consider the `.nl` top-level domain. Next to 4 unicast deployments, SIDN also operates `.nl` with a multi-site anycast deployment. Each site collects statistics on incoming queries at each server. and two servers for `.nl` are located very close to Root DNS servers. Figure 15 shows query counts taken at these severs. We see that these

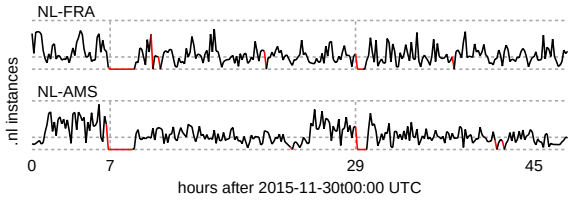

Figure 15: Normalized number of queries for `.nl`, measured at the servers in 10 minutes bins.

sites serve no queries during most of both events. These queries were picked up by other anycast sites for `.nl`, but this data shows collateral damage on a TLD from the attack on the root.

4. RELATED WORK

Distributed Denial-of-Service attacks is a broad area of study and it has been addressed from many different angles in the past years. Studies have shown that DDoS attacks are effective [48].

DDoS attacks are common and growing: Arbor has documented their increasing use and growth in size [2, 3], and we have seen DDoS attacks currently reaching almost 350 Gb/s [50]. Very large attacks often use different protocols to amplification basic attack traffic [40, 45, 16]. Yet DDoS-for-hire ("Booter" services) are easily available for purchase on the gray market—for only a few U.S. dollars, Gb/s attacks can be ordered on demand [42].

Some approaches have been proposed to mitigate amplification [46, 25], spoofing [20] or collateral damage [14]. The continued and growing attacks show that mitigation has been incomplete and spoofing is still widespread [7].

Many studies have looked at the DNS root server system, considering performance [9, 22, 10, 27, 15, 5, 43, 29, 12, 26, 19, 28, 6], client-server affinity [43, 8], and effects of routing on anycast [4, 11], as well a proposal to improve anycast performance in CDNs [21]. We draw on prior measurement approaches, particularly the use of CHAOS queries to identify anycast catchments [19].

Closest to our work are prior analysis of the Nov. 30 events [39, 49, 47, 55]. These reports are important, but were high level [39, 55] or reported only on specific letters [49, 47, 55].

To the best of our knowledge, our paper is the first to combine multiple sources of measurement data to assess how a DDoS attack affects the several layers of the anycast deployment of DNS Root service. In addition, we are aware of no prior public studies on diverse anycast infrastructure operating under stress, including at the site and server level and its consequences on other services (collateral damage).

5. CONCLUSIONS

This paper provides the first evaluation of anycast services under DDoS. Our work evaluates the Nov. 30 and Dec. 1, 2015 events on the Root DNS, evaluating the effects of those events on 10 different architectures, with most analysis based on publicly available data. Our analysis shows different behaviors across different letters (each a separate anycast services), at different sites of each letter, and at servers inside some sites. We identify the role of different policies at overloaded anycast sites: the choice to absorb attack traffic to protect other sites, or to withdraw service in hope that other sites can cover. We believe overall DNS service was robust to this attack, due to caching and the availability of multiple letters for service. However, we show that large attacks can overwhelm some sites of some letters. In addition, we show evidence that high traffic on one service can result in collateral damage to other services in the same datacenter.

Our study shows the importance of anycast design for critical infrastructure, and opens the door for future study in alternative policies that may improve resilience.

Acknowledgments

The authors would like to thank Arjen Zonneveld, Jelte Jansen, Duane Wessels, Ray Bellis, Romeo Zwart, Colin Petrie, Matt Weinberg and Piet Barber for their valuable comments on paper drafts.

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

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

Ricardo de O. Schmidt and Wouter de Vries' work is sponsored by SAND and DAS (http://www.das-project.nl) projects.

John Heidemann and Lan Wei's work is partially sponsored by the U.S. Dept. of Homeland Security (DHS) Science and Technology Directorate, HSARPA, Cyber Security Division, via SPAWAR Systems Center Pacific under Contract No. N66001-13-C-3001, and via BAA 11-01-RIKA and Air Force Research Laboratory, Information Directorate under agreement numbers FA8750-12-2-0344 and FA8750-15-2-0224. The U.S. Government is authorized to make reprints for Governmental purposes notwithstanding any copyright. The views contained in herein are those of the authors and do not necessarily represent those of DHS or the U.S. Government.

6. REFERENCES

[1] Abley, J., and Lindqvist, K. Operation of anycast services. RFC 4786, Internet Request For Comments, Dec. 2006. (also Internet BCP-126).

[2] Arbor Networks. Worldwide infrastructure security report. Tech. Rep. 2012 Volume VIII, Arbor Networks, Sept. 2012.

[3] Arbor Networks. Worldwide infrastructure security report. Tech. Rep. Volume IX, Arbor Networks, Jan. 2014.

[4] Ballani, H., and Francis, P. Towards a Global IP Anycast Service. In *Proceedings of the ACM SIGCOMM* (2007), pp. 301–312.

[5] Ballani, H., Francis, P., and Ratnasamy, S. A Measuremnet-based Deployment Proposal for IP Anycast. In *Proceedings of the ACM Internet Measurement Conference* (2006), IMC, pp. 231–244.

[6] Bellis, R. Researching F-root Anycast Placement Using RIPE Atlas. https://labs.ripe.net/Members/ray_bellis/researching-f-root-anycast-placement-using-ripe-atlas, 2015.

[7] Beverly, R., Berger, A., Hyun, Y., and k claffy. Understanding the efficacy of deployed internet source address validation filtering. In *Proceedings of the ACM Internet Measurement Conference* (Chicago, Illinois, USA, Nov. 2009), ACM, pp. 356–369.

[8] Boothe, P., and Bush, R. Anycast Measurements Used to Highlight Routing Instabilities. NANOG 34, 2005.

[9] Brownlee, N., kc claffy, and Nemeth, E. DNS Root/gTLD Performance Measurement. In *Proceedings of the USENIX LISA conference* (2001), pp. 241–255.

[10] Brownlee, N., and Ziedins, I. Response Time Distributions for Global Name Servers. In *Proceedings of the International conference on Passive and Active Measurements* (2002), PAM.

[11] Bush, R. DNS Anycast Stability: Some Initial Results. CAIDA/WIDE Workshop, 2005.

[12] Castro, S., Wessels, D., Fomenkov, M., and Claffy, K. A Day at the Root of the Internet. *ACM Computer Communication Review 38*, 5 (2008), 41–46.

[13] Chirgwin, R. Linode: Back at last after ten days of hell. The Register, http://www.theregister.co.uk/2016/01/04/linode_back_at_last_after_ten_days_of_hell/, Jan. 2016.

[14] Chou, J. C.-Y., Lin, B., Sen, S., and Spatscheck, O. Proactive surge protection: a defense mechanism for bandwidth-based attacks. *IEEE/ACM Transactions on Networking (TON) 17*, 6 (2009), 1711–1723.

[15] Colitti, L. Effect of anycast on K-root. 1st DNS-OARC Workshop, 2005.

[16] Czyz, J., Kallitsis, M., Gharaibeh, M., Papadopoulos, C., Bailey, M., and Karir, M. Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks. In *Proceedings of the 2014 Conference on Internet Measurement Conference* (New York, NY, USA, 2014), IMC '14, ACM, pp. 435–448.

[17] Eastlake, D. E., and Andrews, M. Domain Name System (DNS) cookies. Work in progress (Internet draft draft-ietf-dnsop-cookies-10.txt), Apr. 2016.

[18] Elz, R., Bush, R., Bradner, S., and Patton, M. Selection and operation of secondary DNS servers. RFC 2182, Internet Request For Comments, July 1997. (also Internet BCP-16).

[19] Fan, X., Heidemann, J., and Govindan, R. Evaluating anycast in the Domain Name System. In *Proceedings of the IEEE Infocom* (Turin, Italy, Apr. 2013), IEEE, pp. 1681–1689.

[20] Ferguson, P., and Senie, D. Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing. RFC 2267, Internet Request For Comments, May 2000. also BCP-38.

[21] Flavel, A., Mani, P., Maltz, D., Holt, N., Liu, J., Chen, Y., and Surmachev, O. Fastroute: A scalable load-aware anycast routing architecture for modern CDNs. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)* (2015), pp. 381–394.

[22] Fomenkov, M., kc claffy, Huffaker, B., and Moore, D. Macroscopic Internet Topology and Performance Measurements From the DNS Root Name Servers. In *Proceedings of the USENIX LISA conference* (2001), pp. 231–240.

[23] Gettys, J., and Nichols, K. Bufferbloat: dark buffers in the Internet. *Communications of the ACM 55*, 1 (Jan. 2012), 57–65.

[24] Holterbach, T., Pelsser, C., Bush, R., and Vanbever, L. Quantifying interference between measurements on the RIPE Atlas platform. In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference* (2015), ACM, pp. 437–443.

[25] Kührer, M., Hupperich, T., Rossow, C., and Holz, T. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *23rd USENIX Security Symposium* (2014), pp. 111–125.

[26] Lee, B.-S., Tan, Y. S., Sekiya, Y., Narishige, A., and Date, S. Availability and Effectiveness of Root DNS servers: A long term study. In *Proceedings of the IEEE Network Operations and Management Symposium* (2010), NOMS, pp. 862–865.

[27] Lee, T., Huffaker, B., Fomenkov, M., and kc claffy. On the problem of optimzation of DNS root servers' placement. In *Proceedings of the International conference on Passive and Active Measurements* (2003), PAM.

[28] Liang, J., Jiang, J., Duan, H., Li, K., and Wu, J. Measuring Query Latency of Top Level DNS Servers. In *Proceedings of the 14th International conference on Passive and Active Measurements* (2013), PAM, pp. 145–154.

[29] Liu, Z., Huffaker, B., Fomenkov, M., Brownlee, N., and kc claffy. Two Days in the Life of the DNS Anycast Root Servers. In *Proceedings of the 8th International conference on Passive and Active Measurements* (2007), PAM, pp. 125–134.

[30] Mockapetris, P. Domain names - implementation and specification. RFC 1035 (INTERNET STANDARD), Nov. 1987.

[31] Operators, H. Personal communication, Apr. 2016.

[32] 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.

[33] ProtonMail. Guide to DDoS protection. https://protonmail.com/blog/ddos-protection-guide/, 2015.

[34] RIPE NCC. Dnsmon. web site https://atlas.ripe.net/dnsmon/, 2015.

[35] RIPE NCC. Ripe atlas root server data. web site https://atlas.ripe.net/measurements/ID, 2015. ID is the per-root-letter experiment ID: A: 10309, B: 10310, C: 10311, D: 10312, E: 10313, F:10304, G: 10314, H: 10315, I: 10305, J: 10316, K: 10301, L: 10308, M: 10306.

[36] RIPE NCC Staff. RIPE Atlas: A global Internet measurement network. *The Internet Protocol Journal 18*, 3 (Sept. 2015), 2–26.

[37] RIPE Network Coordination Centre. RIPE Atlas. https://atlas.ripe.net, 2015.

[38] Root Operators. *www.root-servers.org website*, Apr. 2016.

[39] Root Server Operators. Events of 2015-11-30 http://root-servers.org/news/events-of-20151130.txt, Nov. 2015.

[40] Rossow, C. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Network and Distributed System Security (NDSS) Symposium* (2014).

[41] RSSAC. RSSAC advisory on measurements of the root server system. Tech. Rep. RSSAC002, ICANN, Nov. 2014.

[42] Santanna, J. J., van Rijswijk-Deij, R., Hofstede, R., Sperotto, A., Wierbosch, M., Zambenedetti Granville, L., and Pras, A. Booters-An analysis of DDoS-as-a-service attacks. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on* (2015), IEEE, pp. 243–251.

[43] Sarat, S., Pappas, V., and Terzis, A. On the use of Anycast in DNS. In *Proceedings of the 15th International Conference on Computer Communications and Networks* (2006), ICCCN, pp. 71–78.

[44] Shaikh, A., Kalampoukas, L., Dube, R., and Varma, A. Routing stability in congested networks: Experimentation and analysis. In *Proceedings of the ACM SIGCOMM Conference* (Stockholm, Sweeden, Aug. 2000), ACM, pp. 163–174.

[45] van Rijswijk-Deij, R., Sperotto, A., and Pras, A. DNSSEC and Its Potential for DDoS Attacks: a comprehensive measurement study. In *Internet Measurement Conference (IMC)* (2014), pp. 449–460.

[46] Vixie, P. Response Rate Limiting in the Domain Name System (DNS RRL). blog post http://www.redbarn.org/dns/ratelimits, June 2012.

[47] Weinberg, M., and Wessels, D. Review and analysis of anonmalous traffic to A-Root and J-Root (Nov/Dec 2015). In *24th DNS-OARC Workshop* (Buenos Aires, Argentina, Apr. 2016). (presentation).

[48] Welzel, A., Rossow, C., and Bos, H. On Measuring the Impact of DDoS Botnets. In *7th European Workshop on System Security (EuroSec)* (2014).

[49] Wessels, D. Verisign's perspective on recent root server attacks. CircleID blog post http://www.circleid.com/posts/20151215_verisign_perspective_on_recent_root_server_attacks/, Dec. 15 2015.

[50] Whittaker, Z. 'Largest' denial-of-service attack hit Asian datacenter this year. http://www.zdnet.com/article/largest-denial-of-service-attack-ever-detected-hit-asian-datacenter/, April 2015.

[51] Woolf, S., and Conrad, D. Requirements for a mechanism identifying a name server instance. RFC 4892, Internet Request For Comments, June 2007.

[52] Yan, H., Oliveira, R., Burnett, K., Matthews, D., Zhang, L., and Massey, D. BGPmon: A real-time, scalable, extensible monitoring system. In *Proceedings of the IEEE Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH)* (Washington, DC, USA, Mar. 2009), IEEE, pp. 212–223.

[53] 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.

[54] Zhu, L., Hu, Z., Heidemann, J., Wessels, D., Mankin, A., and Somaiya, N. Connection-oriented DNS to improve privacy and security. In *Proceedings of the 36th IEEE Symposium on Security and Privacy* (San Jose, Californa, USA, May 2015), IEEE, pp. 171–186.

[55] Zwart, R., and Buddhdev, A. Report: K-root on 30 November and 1 December 2015. RIPE Labs blog https://labs.ripe.net/Members/romeo_zwart/report-on-the-traffic-load-event-at-k-root-on-2015-11-30, Feb. 2015.