# Investigating the impact of DDoS attacks on DNS infrastructure

Raffaele Sommese
University of Twente

KC Claffy
CAIDA/UC San Diego

Roland van Rijswijk-Deij
University of Twente

Arnab Chattopadhyay
University of Twente

Alberto Dainotti
Georgia Institute of Technology

Anna Sperotto
University of Twente

Mattijs Jonker
University of Twente

## ABSTRACT

Denial of Service (DDoS) attacks both abuse and target core Internet infrastructures and services, including the Domain Name System (DNS). To characterize recent DDoS attacks against authoritative DNS infrastructure, we join two existing data sets – DoS activity inferred from a sizable darknet, and contemporaneous DNS measurement data – for a 17-month period (Nov. 20 - Mar. 22). Our measurements reveal evidence that millions of domains (up to 5% of the DNS namespace) experienced a DoS attack during our observation window. Most attacks did not substantially harm DNS performance, but in some cases we saw 100-fold increases in DNS resolution time, or complete unreachability. Our measurements captured a devastating attack against a large provider in the Netherlands (TransIP), and attacks against Russian infrastructure. Our data corroborates the value of known best practices to improve DNS resilience to attacks, including the use of anycast and topological redundancy in nameserver infrastructure. We discuss the strengths and weaknesses of our data sets for DDoS tracking and impact on the DNS, and promising next steps to improve our understanding of the evolving DDoS ecosystem.

## CCS CONCEPTS

• **Networks → Naming and addressing**; *Network measurement*;
• **Security and privacy → Denial-of-service attacks**;

**ACM Reference Format:**
Raffaele Sommese, KC Claffy, Roland van Rijswijk-Deij, Arnab Chattopadhyay, Alberto Dainotti, Anna Sperotto, and Mattijs Jonker. 2022. Investigating the impact of DDoS attacks on DNS infrastructure. In *ACM Internet Measurement Conference (IMC '22), October 25–27, 2022, Nice, France.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3517745.3561458

## 1 INTRODUCTION

Distributed Denial of Service attacks are one of the most critical threats on the modern-day Internet. They are cheap, effective, and keep growing in intensity [13, 14, 20]. DDoS attacks that impact the Domain Name System (DNS) are of particular concern, since DNS serves as a support infrastructure for most applications, content distribution platforms, and many security services [15]. If you can stop the DNS, you can effectively stop most Internet communication.

One challenge in preventing such attacks is that they are often indistinguishable from regular DNS traffic, and mitigations may introduce new problems. For example, one mitigation approach is *ingress traffic rate limiting*, which affects not only malicious but also legitimate traffic. The rise of DDoS attacks has paved the way to a new market for DDoS protection systems, i.e. appliances and services aimed at stopping malicious traffic from hindering a certain service. However, cost or privacy constraints may limit the usability of such services. Moreover, they introduce a single point of failure by aggregating traffic toward a single entity.

The persistent DDoS problem triggers questions regarding how pervasive DDoS attacks against critical infrastructure actually are, and what impact they have. Attackers generally know they are launching an attack (although not necessarily how successful it is), and a victim of a successful attack is generally aware of it due to service impairment, but may not want to publicize that fact. But independent study of DoS attacks at scale is a long-standing challenge. A third party has to contend with discerning an attack from myriad root causes of service impairment on the global Internet. Heavily capitalized players can put significant resources into monitoring millions of IP addresses in network traffic [28], but these approaches are difficult to scale, and not within reach of academic research efforts.

In this work we develop a scalable method to map DDoS attacks targeting or affecting DNS infrastructure. We use two unique macroscopic data sets to develop this mapping: the UCSD Network Telescope, which collects backscatter traffic from ongoing DDoS attacks against IPv4 address space; and the OpenINTEL measurement project, which performs daily DNS queries of over 60% of registered domains, allowing detection of substantial changes in DNS query latency or reachability to authoritative nameservers over time. Resolution times experienced by OpenINTEL during attacks indicate their impact on the DNS; network telescope traffic allows partial inferences of attack timing and intensity.

We expand on the following contributions:

(1) We synthesized two data sets that capture global IPv4 behavior to discover evidence of attacks against tens of millions of domains (≈ 5% of the DNS namespace), although often with negligible performance impact.
(2) We discovered attacks against DNS providers that impaired performance and reachability for millions of domains.


*IMC '22, October 25–27, 2022, Nice, France*

ACM ISBN 978-1-4503-9259-4/22/10.
https://doi.org/10.1145/3517745.3561458

(3) Our data confirms the effectiveness of the use of anycast and diversity in nameserver deployment in providing resiliency against DDoS attacks.
(4) We document corroborating evidence of politically motivated attacks on Russian infrastructure.
(5) We analyze the limitations of our data sets to infer effectiveness of attacks, and propose approaches to overcome them in our pursuit of more accurate characterization the DoS ecosystem.

Our study illustrates the value of combining longitudinal datasets in extracting cybersecurity-related insights into Internet evolution, in this case regarding the observable harms of DDoS attacks to performance and availability of critical services.

## 2 BACKGROUND

### 2.1 DDoS Attacks

Distributed Denial of Service (DDoS) attacks are a notorious type of cyberattack. While conceptually simple, DDoS attacks can be highly effective in disrupting networks and denying users access to on-line services. Attackers are known to misuse core Internet infrastructure to bring about attacks, as well as target it. With society ever-increasingly relying on the Internet as its communications fabric, the persistent threat that DDoS poses against Internet stability and reliability is nothing short of grievous.

By and large, attacks can be classified as *volumetric* or *semantic*. The prior involves using sheer network traffic volumes (e.g., high packet rate and/or byte magnitude). The latter involves abusing specific weaknesses (e.g, in L7 protocols) without relying on a high rates per se. We distinguish between three categories of attacks for the purpose of this background. *Unspoofed* attacks involve sending network traffic directly from the attacking infrastructure (e.g., IoT botnet) to the victim host, without application of source IP address spoofing. *Reflected* (or indirect) attacks involve specific source IP address spoofing, to dupe so-called reflectors (e.g., open DNS resolvers) to send traffic to the victim host, in response to requests purportedly coming from the victim host. Finally, *randomly spoofed* attacks involve randomly (and often uniformly) spoofing the source IP address, in attempt to conceal the attacking infrastructure.

Obtaining data on DDOS attacks is non-trivial. Inferring attacker behavior in the case of reflected attacks requires complex honeypot reflectors to mimic frequently used sources. Detection of spoofed attacks requires access to a large source of backscatter traffic, i.e., a large darknet. Finally, detection of unspoofed attacks requires the collaboration of victims and/or network providers, who are not generally sharing such data. The challenges with data access limit the ability of researchers to characterize the evolution of DDoS attacks on the Internet. The two longitudinal data sets available to us allow a focus on randomly and uniformly spoofed attacks. Sizable attacks of this type will use many spoofed IP addresses and thus appear as sourced from a wide range of networks, captured by both of our data sets.

### 2.2 DNS and IP Anycast

The domain name system is the Internet's phonebook. Its primary task is to map names to IP addresses. The DNS comes as a distributed and decentralized database. It was designed with reliability in mind. For example, RFC 1034 [21] requires every zone the be available on at least two authoritative nameservers. RFC 2182 [29] further recognizes that diversity in terms of topological and geographical placement of redundant servers increases reliability. Ironically, the number of root server IP addresses is capped at thirteen. In the early 2000s, however, operators of DNS root servers started distributing replicas of these servers around the world, for which they rely on IP anycast. IP anycast leverages the border gateway protocol to allow multiple server instances to use the same IP address. This was a successful endeavor that did not go unnoticed. Anycast has since been adopted in numerous other services to add resilience and is being used at other levels of the DNS hierarchy (e.g., public resolvers such as Quad9 and top-level domain authoritatives). Anycast deployment however requires specific knowledge and routing resources. While it is a great way to add resilience for critical infrastructure, arguably it may be superfluous for others. Finally, the DNS comes with caching mechanisms for performance, and to reduce the likeliness of resolution failure in case of intermittent connectivity issues. The strong rise in use of content delivery networks, however, reduces the effectiveness of caching, as CDNs typically configure lower cache lifetimes (i.e., time-to-live values) to aid with DNS-based load-balancing.

## 3 DATASET

We join two primary, long-standing datasets for study. To get indicators of DoS attacks against IPv4 address space, we use inferences made from UCSD Network Telescope (UCSD-NT) data. To study which DNS authoritative nameservers exhibit performance degradations, we use contemporaneous DNS measurement data from the OpenINTEL project. We also use several ancillary datasets to support our analysis.

### 3.1 DoS Attacks Inferred from Internet Background Radiation (IBR)

The UCSD-NT announces and captures traffic destined to two globally routed networks – a /9 and /10 address block, covering approximately 1/341 of the total IPv4 address space. The collected traffic is referred to as *Internet Background Radiation (IBR)*, a significant component of which is *backscatter*, including packets that are sent in response to randomly spoofed DDoS attacks. CAIDA curates the raw data to create a *Randomly and Uniformly Spoofed Denial of Service (RSDoS) attacks* feed that consists of a 5-minute tumbling window of aggregated statistics of response packets sent by victims of RSDoS attacks [9]. We use this data feed to establish a lower bound of DoS attacks against specific IP addresses. In addition to a timestamp of each 5-minute window, this data set includes several fields that we use to characterize attacks: the number of /16 subnets in the telescope that receive packets from the inferred victim in the 5-minute window; protocol, first observed port, and number of unique ports targeted; and peak observed packet rate during the window. The first port indicates which service was under attack in single-port attacks. The RSDoS data contains 4,039,485 inferred attacks for November 2020 to March 2022 (Table 1).

| #Attacks | #IPs | #/24 Prefixes | #ASes |
|---|---|---|---|
| 4,039,485 | 1,022,102 | 404,076 | 25,821 |

**Table 1: RSDoS Dataset: November 2020-March 2022**

### 3.2 OpenINTEL - Active DNS Measurements

OpenINTEL is a large-scale measurement platform that performs daily querying of all domain names registered under many top-level domains (TLDs), including all gTLDs participating in ICANN's Centralized Zone Data Service (CZDS) platform, legacy gTLDs (*e.g.,* `.com`, `.net`, `.org`) and several ccTLDs (*e.g.,* `.at`, `.nl` and `.ru`) [36]. OpenINTEL also measures domain names included in various Top lists. OpenINTEL performs several queries, including NS queries, for each domain name, and stores the round trip time (RTT) to complete the query, along with response status codes (*e.g.,* OK, SERVFAIL, TIMEOUT). OpenINTEL triggers explicit NS queries to deal with parent-child inconsistency, and prefers the authoritative answer [34]. Explicit NS queries trigger a direct response from the targeted authoritative nameservers, providing us the effective RTT to reach them. The query process uses DNS resolver software *unbound* [26] to randomly select an authoritative nameserver for the first query for every registered domain (*i.e.,* excluding caching[1]). This *agnostic* resolver behavior captures the actual resilience mechanisms implemented by DNS operators, but also prevents us from identifying which specific authoritative nameserver responded to each query.

### 3.3 Anycast Census and Additional Datasets

We use quarterly census snapshots of anycast deployment taken from January 2021 until January 2022 [33]. We identify DNS anycast deployments within this snapshot by matching authoritative NS IP /24 subnets with /24 subnets detected in the anycast census, as in [32]. This census provides us a lower bound estimation of anycast deployments. We also leverage CAIDA's prefix-to-AS dataset [8] to map IP addresses to the AS number(s), and CAIDA's AS-to-organization [7] to map AS numbers to organizations. Finally, we use the open resolver scans of Yazdani *et al.* [38] to filter out incidental IPs of open resolvers showing up in the DNS authoritative infrastructure.

## 4 METHODOLOGY

Our methodology consists of four steps: (1) Create an aggregated dataset of the OpenINTEL data; (2) Map IP addresses under attack to nameservers under attack; (3) Extract the list of domains associated with those nameservers; (4) Use the RTT data to infer performance impairment for queries to those domains. Our analysis interval is the 17-month period from November 1, 2020 to March 31, 2022, which lines up with the anycast census data (§3.3).

### 4.1 Extrapolating DNS Performance Metrics

OpenINTEL does not record which authoritative nameserver provided the answer to a query for a specific domain, so we aggregated performance metrics for all IPv4 nameserver IP addresses in common for one or more domains, which we define as its *NSSet*. This aggregation allows us to estimate the impact of nameserver deployment scenarios on resolution performance. Each NSSet contains the IP addresses of the authoritative nameservers as well as their corresponding autonomous system number (ASN), prefix, and country code. For each NSSet, we collect, in a 5-minute interval (*i.e.,* the same granularity of RSDoS attack dataset), the number of domains resolved by OpenINTEL, and the average, minimum, and maximum RTT observed for that interval, and number of errors (*e.g.,* Timeout, SERVFAIL, etc.). This data allows us to define the following metric of the impact of an attack on RTT of queries for a domain, and thus the impact on end users:

$$Impact\_on\_RTT = \frac{Average\ RTT\ (5\ min)}{Average\ RTT\ (Day\ Before)} \quad (1)$$

Significant RTT increases above a baseline are indicative of either an attack causing network congestion or other path impairments. By joining the OpenINTEL data with RSDoS data, we can correlate RTT changes with inferred RSDoS attacks. We evaluated using different time-window metrics as a baseline (*e.g.,* Average RTT (Week/Month Before)) finding similar results. We decided to stick with the previous day metric to minimize errors due to infrastructural changes in the DNS hosting architecture. While averaging RTT may mask outliers, it provided us a stable metrics to identify the impact of DDoS attacks.

### 4.2 Joining datasets

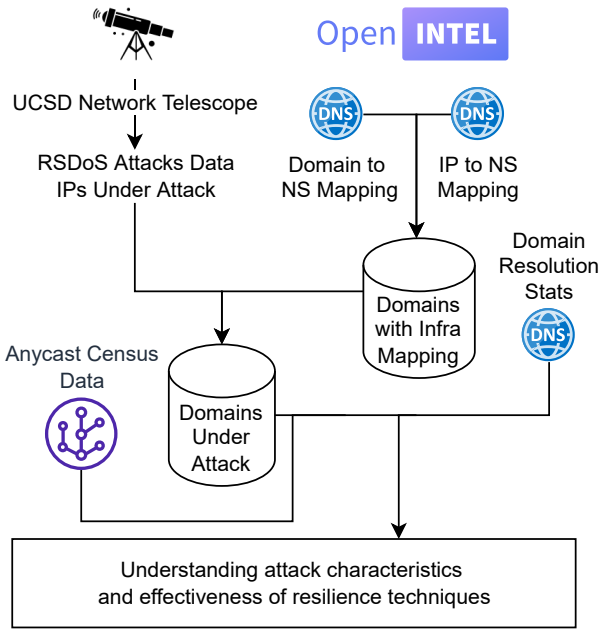


**Figure 1: Data analysis pipeline: The RSDoS Feed joined with the OpenINTEL measurement provides information on the impact on DNS infrastructure during DDoS attacks**

[1] Additional queries may leverage cached NS or other records, providing a successful resolution of domains under attack reducing the visibility on the real impact of attacks.

Figure 1 shows how we join the RSDoS-compiled IPs under attack with the list of nameservers successfully queried on the day before the attack. This step maps IPs under attack to nameservers under attack. We join the resulting dataset with the list of domain names those nameservers hosted, as observed on the day before the attack. This step yields the list of domains under attack. By using the list of nameservers on the previous day, we minimize the chance of missing a nameserver that is unreachable due to an attack. We assume daily changes in nameserver infrastructure will not significantly affect our analysis. We join the list of domains under attack with our RTT data for NSSets. We use additional metadata (subsection 3.3) to characterize performance during various attack windows.

### 4.3 Limitations

The following limitations of our data sets constrain our inference capability. First, OpenINTEL's *agnostic* DNS resolution (§3.2) means that we cannot know which authoritative nameserver responded to a query. The random selection of authoritative nameserver means that eventually it should query each nameserver, but it also restricts our ability to discern behavior of (and thus the impact of attacks on) different nameservers for the same domain. While this limitation does not allow us to pinpoint the behavior of single nameservers affected by the attack, it enables inference of how a typical end user would experience DNS resolution. Therefore, we can estimate a realistic worst case scenario of end-user experience in resolving a domain with an empty cache. This empty-cache configuration implies that the TTL value for a specific domain will not impact the resolution performance of OpenINTEL.

Second, OpenINTEL resolves domains using both IPv4 and IPv6 addresses, but the RSDoS data includes only IPv4 addresses. During an attack on IPv4 DNS infrastructure, separate parallel IPv6 infrastructure might be operational, limiting the impact of an attack. On the other hand, often IPv4 and IPv6 services share the same infrastructure and even server [6], in which case our inferences would hold.

Third, the network telescope detects only a specific kind of attack, which uses randomly spoofed IP addresses to launch a volumetric attack. During multi-vector attacks, we have limited visibility of overall attack duration and intensity. We also have no visibility into reflected and unspoofed attacks. As a relevant data point, in Jonker *et al.* [13] compared two data sets of inferred attacks over two years, finding 60% of attacks as randomly spoofed (observed in RSDoS data), and 40% as reflected attacks (observed in the AmpPot data).

Finally, the single vantage point from which OpenINTEL queries in a highly complex Internet topology limits the precision of our visibility of the performance impact of attacks, especially in case of anycast deployments where catchment can mask ongoing attacks in specific geographic regions.

*4.3.1 Reactive measurement.* To mitigate these limitations, we have built a reactive measurement platform that iteratively targets the full list of authoritative nameservers when resolving a domain name. Every time an RSDoS feed reports an attack, our platform joins the list of IPs under attack with the list of nameservers provided by OpenINTEL and the registered domain names that delegate authority to said servers. For every attack, we trigger probes of 50 related domain names every 5 minutes during the attack and in the 24 hours after (to characterize the post-attack baseline behavior). We choose this threshold to avoid additional burden on infrastructure already overloaded by attacks. Moreover, we spread our 50 measurements over the entire 5-minute window. We launched these measurements operationally in January 2022, so could not use them for our longitudinal analysis, but we did leverage them to study the impact of attacks against Russian infrastructure (subsection 5.2). Although our current infrastructures probes from a single vantage point in the Netherlands, we are in process of additional vantage points to increase visibility of how attacks affect performance and availability in different geographic regions (*e.g.,* due to anycast catchment).

We built our analysis pipeline using Kafka [18], Spark Structured Streaming [39], Apache Flume, and Grafana to display results. We use this pipeline to trigger reactive measurements with a maximum delay of 10 minutes after the start of an attack. In the future we can use this platform to perform near real-time characterization of DDoS attacks on DNS infrastructure.

## 5 RESULTS: CASE STUDIES

### 5.1 Large European hosting provider

| Target Nameserver | | A | B | C |
|---|---|---|---|---|
| December 2020 Attack | Observed Packer Rate (PPM) | 21.8K | 3.8K | 2.9K |
| | Inferred Traffic Volume | 1.4 Gbps | 247 Mbps | 188 Mbps |
| | Attacker IP Count | 5.79M | 1.57M | 1.33M |
| March 2021 Attack | Observed Packer Rate (PPM) | 125K | 123K | 13K |
| | Inferred Traffic Volume | 8 Gbps | 7.8 Gbps | 845 Mbps |
| | Attacker IP Count | 7M | 6.19M | 823K |

**Table 2: Attack metrics for two DDoS attacks on TransIP. The first attack targeted nameserver A more intensely; the second targeted all three similarly.**

Our first case study exemplifies how DDoS attacks can impact large providers, severely degrading DNS performance for end users. We focus on two attacks against TransIP, a large European DNS and hosting provider. Both attacks were reported [12, 35] and acknowledged by TransIP [27]. At the time of the two attacks (December 2020 and March 2021) TransIP was responsible for ≈ 8% of `.nl` domains, potentially affecting millions of users. By joining the two data sets, we infer that these attacks potentially affected ≈ 776K domain names, two-thirds of which (≈510K) were `.nl` domains. At the time of the attacks (and still in May 2022), TransIP used three unicast IPs for nameservers for the domain names they hosted, all of which appeared as RSDoS attack targets (A, B, and C in Table 2).

In December 2020, the network telescope data shows evidence of RSDoS attack activity from 2020-11-30 at 22:00 to 2020-12-01 12:30 (UTC). We estimate an attack rate of 124Kpps (21.8K packets

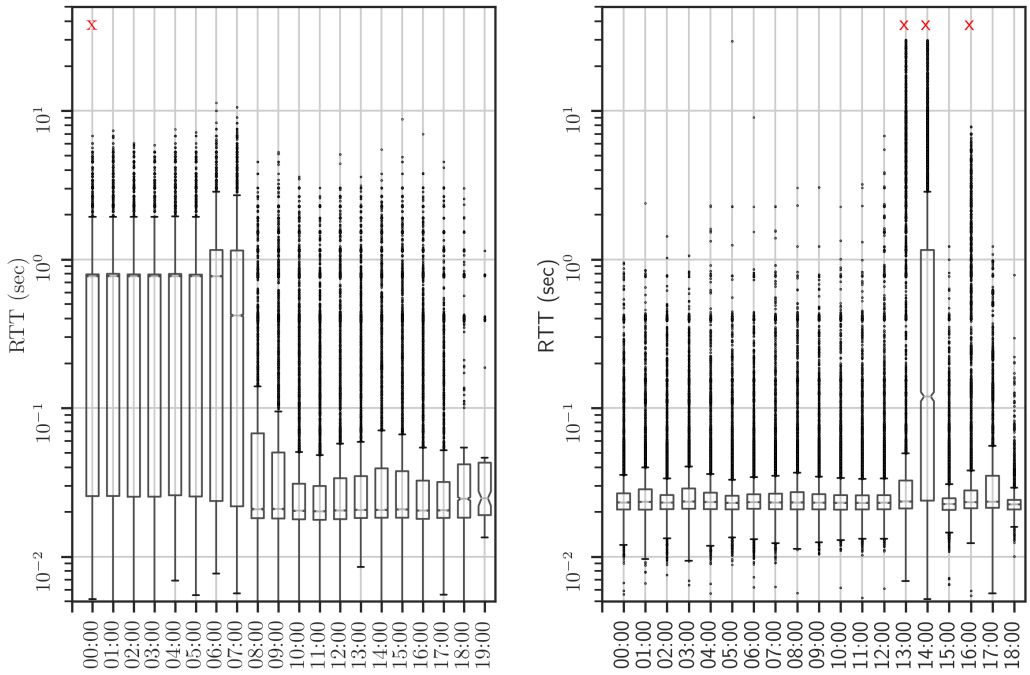


(a) December 2020 (b) March 2021

**Figure 2: RTT variations in DDoS attacks on TransIP. The attack hours are marked with a red cross. Effects of the December attack persisted for hours after the RSDoS-inferred end of the attack. The March attack induced larger RTT impairments.**

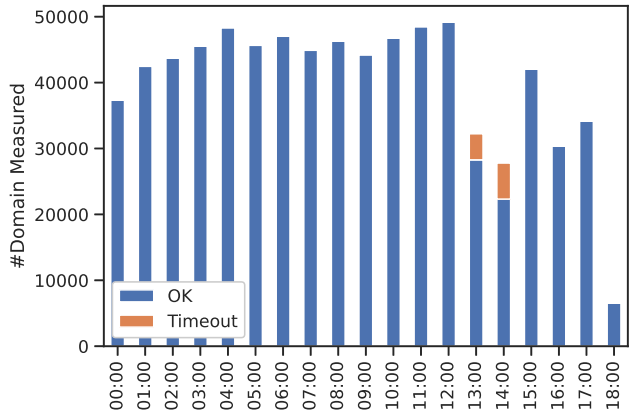


**Figure 3: Timeout errors during the March 2021 attack on TransIP reached 20% of observed domains, leading to resolution failures for end users**

per minute at the telescope[2]) against nameserver A. Nameservers B and C seem to have experienced lower-intensity attacks (Table 2). The lower intensity inferred for B and C suggests low impact on overall DNS operations, but OpenINTEL measured a 10X increase in DNS resolution time, indicating significant impairment (recall that OpenINTEL randomizes which NS to query for each registered domain, thus over 770K domains it is overwhelmingly likely to send a similar number of queries to all three nameservers). The performance impairment ended on 2020-12-01 at 08:00, 8 hours after the RSDoS-inferred end of the attack. One explanation for this behavior is that the attackers moved to a different kind of DDoS attack not observable by the network telescope. Another explanation is the need for human intervention to restore DNS service quality.

In their report, TransIP stated that the March 2021 attack was more intense [27]. Consistently, the telescope observed a peak packet rate 6X greater than for the December 2020 attack and, as shown in Figure 3, ≈ 20% of OpenINTEL queries timed out and failed to resolve (compared to a negligible fraction in December). The March attack more likely impacted end users because it induced complete unreachability of domain names. Nevertheless, differently from the December attack, the time frame in which we observed an impact matched the interval inferred through the telescope data (Figure 2). This difference might be associated with TransIP's reported use of a DDoS protection/scrubbing mechanism [27]. Since OpenINTEL observed no evidence of nameserver changes during the attack, we speculate that the scrubbing service might have been deployed at the IP level.

*5.1.1 Resilience techniques adopted by TransIP.* This case study shows that even with traffic scrubbing, DDoS attacks can affect resolution for hundreds of thousands of domains. More strategic

[2] 21.8kppm × 341 / 60s = 124K pps.

deployment of DNS infrastructure would have improved its resiliency. TransIP served the registered domain using three unicast authoritative nameservers, on three different subnets, in two separate geographic locations (Amsterdam and Eindhoven), behind a single ASN. While hosting these nameservers behind different subnets increased resilience, the lack of anycast deployment left these domains dependent on three physical servers and (at most three) network links. Moreover, hosting these domains within a single ASN means they relied exclusively on a single company's infrastructure. Using a more diversified infrastructure, by using anycast and/or third-party hosting providers, would have further mitigated the effects of these attacks. Finally, our analysis shows that ≈ 27% (203,217) of the domains hosted by TransIP relied on third-party hosting for their web content. These domains likely felt the December attack less, i.e., simply experienced slower DNS resolution but during the March attack they likely became entirely unreachable due to DNS resolution failures, despite having a third party operating their web site.

## 5.2 Attacks on Russian Assets in 2022

The TransIP case illustrates a type of attack against commercial infrastructure, whereas in this second example we discuss attacks targeting infrastructure hosting specific domains and likely motivated by political reasons. Specifically, we focus on several attacks targeting Russian government web sites in March 2022, shortly after Russia's invasion of Ukraine.

*5.2.1 Russian Ministry of Defense.* The first attack targeted *mil.ru*, the domain of the Ministry of Defense of the Russian Federation. Three nameservers on the same /24 subnet were authoritative for both the international and the Cyrillic IDN name of mil.ru and for several subdomains.

These three nameservers were under attack for 8 consecutive days, March 11-18, according to RSDoS inferences from the network telescope data. The telescope detected a modest-intensity attack, although newspapers reported a severe attack on the network infrastructure of mil.ru and other government web sites [5, 30]. Newspapers also reported the geo-fencing of mil.ru in response to the attack, allowing connections only from within the Russian state. OpenINTEL completely failed to resolve mil.ru (and the related Cyrillic domain) for most of the attack period (from March 12 to 16, inclusive), whereas our reactive measurements (§4.3.1) found the domain unresolvable for the entire duration of the attack, with none of the three nameservers responsive.

*5.2.2 RDZ Railways.* Another interesting case study is related to RDZ railways. RSDoS data indicates an attack from 15:30 to 20:45 on March 8, 2022. Our reactive measurement system launched queries to resolve the domain in the 24 hours following the start of the attack, and found the domain became intermittently responsive at 06:00 on the next day. We also found evidence that this attack was co-coordinated via a Telegram channel named *IT ARMY of Ukraine* (Figure 4). A message on the channel at 15:43 provided the IP address of the 3 RDZ railway DNS nameservers, asking for assistance to crowdsource an attack on port 53/UDP, 12 minutes after RSDoS-inferred start time of the first attack.

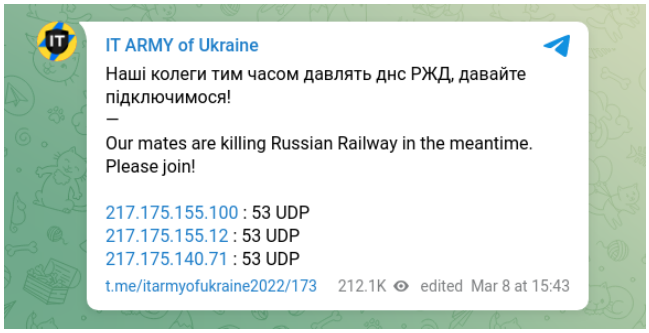


**Figure 4: Telegram coordination of DDoS attack. We manually inspected the messages to find evidence of correlation.**

*5.2.3 Resilience techniques adopted by Russian Infrastructures.* The attack on `mil.ru` is a textbook illustration of poor resilience in DNS infrastructure. The three nameservers were unicast, hosted behind the same ASN/company, and even on the same /24 subnet. This lack of resilience contributed to the attack's success, apparently forcing the Russian government to geofence the entire network to protect DNS service. Moreover, being hosted on the same /24, the three nameservers (and other services on the same subnet, including the `mil.ru` web site) likely shared upstream network resources. This network bottleneck implies that a single intense volumetric attack targeting a service on the network can affect all services hosted on the network. The RDZ railways domain had a slightly more resilient deployment. The three nameservers were hosted on two separate /24 subnets, but still used unicast and a single ASN. However, as in the mil.ru case, the attacker targeted all three nameservers, and simple prefix diversity was not sufficient to withstand the attack.

# 6 LONGITUDINAL ATTACKS ANALYSIS

Although the case studies illustrate the value of joining these two data sets together to corroborate known attacks, our ultimate goal is to use this method to identify and track the prevalence and scope of unreported attacks against global DNS infrastructure in the wild. We used data from November 2020 to March 2022 to identify all RSDoS-inferred attacks against DNS infrastructure, either directly targeting nameserver IPs or targeting /24s that host nameservers.

## 6.1 Overview of Attacks in 2020-2022

Table 3 shows that attacks on the DNS infrastructure are between the 0.57% and 2.12% of total attacks detected by the telescope, spanning ≈ 1 − 2% of the total affected IPs. Although this is a small percentage of the total number of attacks, the IP addresses may be nameservers that host millions of domains. We focused on attacks directly targeting nameserver IPs (rather than the containing subnet or announced prefix). Figure 5 shows the monthly counts of *potentially affected* domains, i.e., one of its nameservers was under attack. On average, 10-100 domains were potentially affected by attacks, although much larger numbers are common. We also estimated the attack's intensity and the handling capacity of the target infrastructure. We identified 8 peaks of potentially 10 million domains affected – a series of attacks trying to target around ∼ 4% of the global DNS infrastructure measured by OpenINTEL. These

| Year | Month | #DNS Attacks | #Other Attacks | Total Attacks | DNS IPs | Other IPs | Total (Unique) IPs |
|---|---|---|---|---|---|---|---|
| 2020 | 11 | 2,550 (1.63%) | 156,884 (98.37%) | 159,434 | 798 (1.64%) | 47,839 (98.36%) | 48,637 |
| | 12 | 3,876 (1.08%) | 356,042 (98.92%) | 359,918 | 1,070 (0.94%) | 113,354 (99.06%) | 114,424 |
| 2021 | 1 | 2,927 (1.68%) | 171,089 (98.32%) | 174,016 | 930 (1.43%) | 63,971 (98.57%) | 64,901 |
| | 2 | 2,873 (1.98%) | 141,949 (98.02%) | 144,822 | 827 (1.52%) | 53,461 (98.48%) | 54,288 |
| | 3 | 3,294 (1.18%) | 276,503 (98.82%) | 279,797 | 929 (0.52%) | 177,514 (99.48%) | 178,443 |
| | 4 | 3,522 (2.12%) | 162,361 (97.88%) | 165,883 | 802 (1.36%) | 58,077 (98.64%) | 58,879 |
| | 5 | 3,973 (1.99%) | 195,540 (98.01%) | 199,513 | 880 (1.19%) | 72,899 (98.81%) | 73,779 |
| | 6 | 2,244 (0.98%) | 227,874 (99.02%) | 230,118 | 821 (0.96%) | 84,294 (99.04%) | 85,115 |
| | 7 | 2,245 (0.66%) | 335,948 (99.34%) | 338,193 | 967 (0.91%) | 105,917 (99.09%) | 106,884 |
| | 8 | 4,473 (1.53%) | 288,369 (98.47%) | 292,842 | 1,055 (1.14%) | 91,517 (98.86%) | 92,572 |
| | 9 | 2,577 (1.05%) | 242,713 (98.95%) | 245,290 | 780 (1.12%) | 68,561 (98.88%) | 69,341 |
| | 10 | 1,968 (0.86%) | 226,124 (99.14%) | 228,092 | 624 (1.25%) | 49,310 (98.75%) | 49,934 |
| | 11 | 2,662 (0.94%) | 281,907 (99.06%) | 284,569 | 835 (1.06%) | 77,942 (98.94%) | 78,777 |
| | 12 | 2,984 (1.35%) | 218,070 (98.65%) | 221,054 | 706 (1.04%) | 67,422 (98.96%) | 68,128 |
| 2022 | 1 | 2,028 (0.86%) | 232,999 (99.14%) | 235,027 | 705 (1.23%) | 56,616 (98.77%) | 57,321 |
| | 2 | 1,368 (0.57%) | 238,407 (99.43%) | 239,775 | 572 (0.88%) | 64,201 (99.12%) | 64,773 |
| | 3 | 3,294 (1.37%) | 237,848 (98.63%) | 241,142 | 669 (0.94%) | 70,778 (99.06%) | 71,447 |
| Total | | 48,858 (1.21%) | 3,990,627 (98.79%) | 4,039,485 | 8,864 (0.87%) | 1,013,238 (99.13%) | 1,022,102 |

**Table 3: Monthly attack activity summary. Attacks toward IPs used as DNS nameservers constituted $\approx 1-2\%$ of the total attacks.**

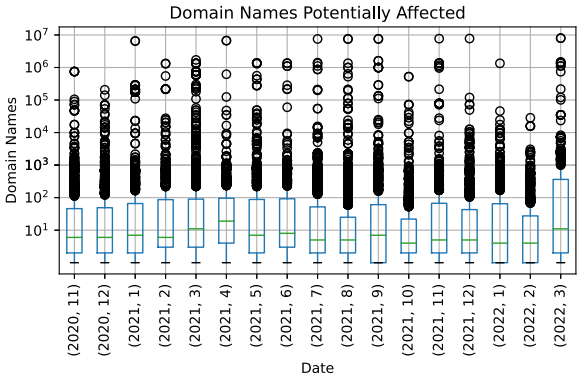


**Figure 5: Registered domains potentially affected by attacks, by month. Some attacks hit deployments serving more than 10 million domain names.**

| ASN | #Attacks | Company | ASN | #Attacks | Company |
|---|---|---|---|---|---|
| 15169 | 7,324 | Google | 16509 | 1,564 | Amazon |
| 46606 | 2,841 | Unified Layer | 8068 | 1,240 | Microsoft |
| 13335 | 2,428 | Cloudflare | 54113 | 1,054 | Fastly |
| 16276 | 2,192 | OVH | 199608 | 894 | Birbir |
| 24940 | 2,172 | Hetzner | 48678 | 562 | Pendc |

**Table 4: Top 10 ASNs attacked from Nov. 20 to Mar. 22. Large DNS hosting companies and cloud providers are the most attacked, usually with negligible effects.**

specific attacks did not substantially harm the performance and operation of these large providers.

We analyzed which companies received the most attacks during our measurement window, finding spikes against Cloudflare and Google DNS infrastructure (Table 4). We analyzed the target IPs for these attacks (Table 5) and found they related to Google's public DNS service (8.8.8.8 and 8.8.4.4) and Cloudflare's Quad1 (1.1.1.1). We see these open resolver nameserver IP addresses in our data due to misconfigured domains pointing their NS records at these IPs. We filtered out such attacks toward open resolvers, since they are not used for authoritative DNS resolution.

We noticed many low-impact attacks against a shared IP address hosted on Unified Layer. After manual inspection, we discovered the IP address has been hosting the web site of an American Youtuber. VirusTotal [37] suggests that the address may have been used in the past for malicious activities. We also found evidence of several attacks against a Russian DNS provider, Beeline, during March 2022. Beeline provided DNS hosting for several Russian banking web sites such as Sberbank, Russian Agricultural Bank, and Eurasian Development Bank. Finally, we found targeted IPs related to Bing and Cloudflare and two private IPs likely related to misconfigured servers leveraged for attacks, belonging to private companies which we omit from our report.

**Key Takeaway:** *Approximately 0.5-2% of RS-DoS attacks observed by the network telescope reached, and perhaps targeted, DNS infrastructure. Some of these attacks hit deployments of 10 Million domain names, with negligible performance impact. Frequent targets included open resolvers, large DNS providers, and hosting companies.*

## 6.2 Targeted Services (Ports)

In our analysis of protocol and port usage by the attacks, we found that 80.7% of attacks to DNS authoritative infrastructure targeted a single port and protocol (Figure 6). Almost 90.4% of these attacks used TCP (mainly TCP SYNs), 8.4% targeted UDP ports, and 1.2%

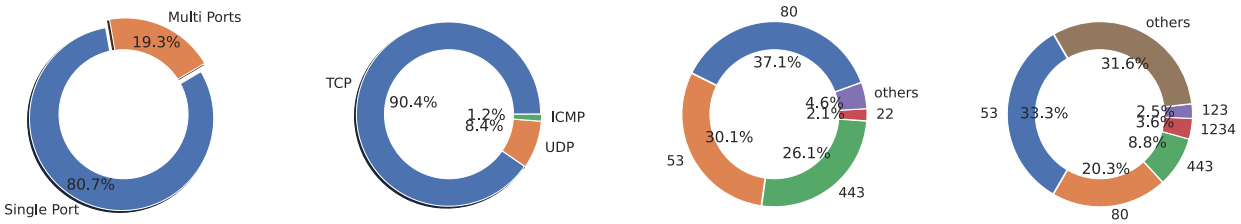


(a) #Targeted Ports (b) Protocol (c) TCP Port (d) UDP Port

**Figure 6: Distribution of protocol and destination ports used by attacks. Most attacks targeted a single port, usually via TCP. The most targeted port was 80 (HTTP), followed by 53 (DNS).**

| IP | #Attacks | Type |
|---|---|---|
| 8.8.4.4 | 2,803 | Google DNS |
| REDACTED | 2,566 | Unified Layer |
| 8.8.8.8 | 2,298 | Google DNS |
| 1.1.1.1 | 1,118 | CloudFlare DNS |
| 204.79.197.200 | 668 | Bing |
| 194.67.7.1 | 481 | Beeline RU |
| 13.107.21.200 | 438 | Bing |
| REDACTED | 400 | Company NAS |
| REDACTED | 346 | Private IP |
| 23.227.38.32 | 273 | Cloudflare |

**Table 5: Top 10 IPs attacked. The presence of open resolver IP address (8.8.4.4, 8.8.8.8, 1.1.1.1) on this list implies that misconfigured domains use them as authoritative NS resolvers. Attacks on such heavily provisioned anycast targets are likely ineffective.**

used ICMP. Historically, DNS was a service provided via UDP. But in the last decade, the introduction of DNSSEC and its need for larger responses led to expanded support for DNS-over-TCP. This fact and the popularity of TCP SYN flood attacks increased the prevalence of TCP-based attacks. The port distribution varies: 37% of TCP-based attacks targeted port 80 (HTTP), and 30% targeted port 53 (DNS); the other significantly used port was 443 (HTTPS). One explanation for the use of HTTP(s) ports against DNS nameserver IPs is the awareness that sometimes the same IP address hosts both DNS and web service. UDP-based attacks exhibited a more varied port distribution, but one-third of UDP attacks targeted port 53.

What stands out is that *the majority of attacks do not target port 53*. This suggests that DNS itself may not be the primary target of these attacks although without knowing an attackers motivations, we cannot be certain of this. Regardless of whether DNS is the target, if the goal of the attacks is to flood the link of the target or to exhaust system resources of the target host, they may still have an impact on DNS resolution. We discuss this further when we consider *successful attacks* in Section 6.3.1.

**Key Takeaway:** *Most attacks towards DNS authoritative nameserver IPs targeted a single port, usually via TCP. The most attacked port was 80 (HTTP), followed by 53 (DNS).*

## 6.3 Performance Impact of Attacks

To assess the performance impact of attacks on DNS infrastructure, we computed a longitudinal 5-minute performance metric based on OpenINTEL's RTT measurements for each NSSet deployment (described in §3.3). To reduce possible sources of noise, we considered only NSSets with at least five domains measured during the attack. Using this constraint, we inferred 12,691 distinct events of attacks to distinct NSSets in the window where OpenINTEL actively measured domains for which the targeted nameserver(s) were authoritative.

*6.3.1 Complete failure in resolution of domain names.* In 99% of these 12,691 attacks, authoritative nameservers continued to provide the answer. However, in 1% of cases we saw domains fail to resolve, with timeout (92%) or SERVFAIL (8%) errors. This result shows that despite most attacks not harming operations, some caused end user failure in resolution (e.g., the TransIP examples discussed in §5.1).

Figure 7 shows the failure rate as a function of the number of domains resolved by OpenINTEL (y-axis); the dot's color represents number of domains (order of magnitude) hosted by the NSSet under attack. Most domains that failed to resolve belonged to small infrastructures. Some attacks induced resolution failures (timeout errors) on large infrastructures hosting > 10K domains. The most effective attack in this size range causing failed resolution for 100% of domains belonged to nic.ru, a Russian registrar. They offer secondary nameservers as a service; those nameservers were attack targets during March 2022.

Most effective attacks occurred against smaller deployments (100-10K domains). A Spanish ISP (Euskaltel) responsible for 1405 domains failed to respond to 83% of queries for its domains during the attack. 99% of domains that failed resolution in this way used unicast nameserver infrastructure.

The impact on end-users in cases of complete resolution failure depends on several factors, mainly related to caching policy. A

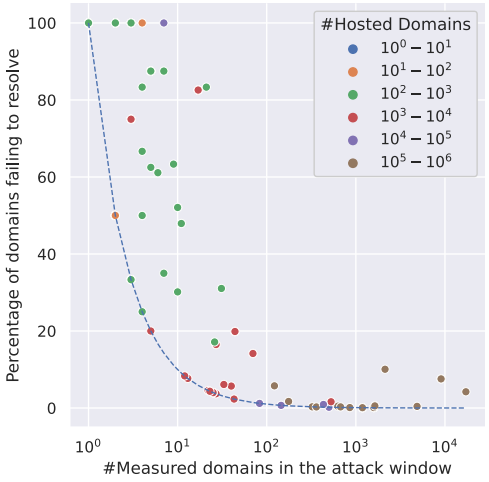


**Figure 7: Percentage of measured domains failing to resolve, colored by number of hosted domain names. The base curve represents a single failure per attack window, everything above this line represents NSsets that experienced failures for multiple domains. Attacks with higher induced failure rates cause complete unreachability for end users.**

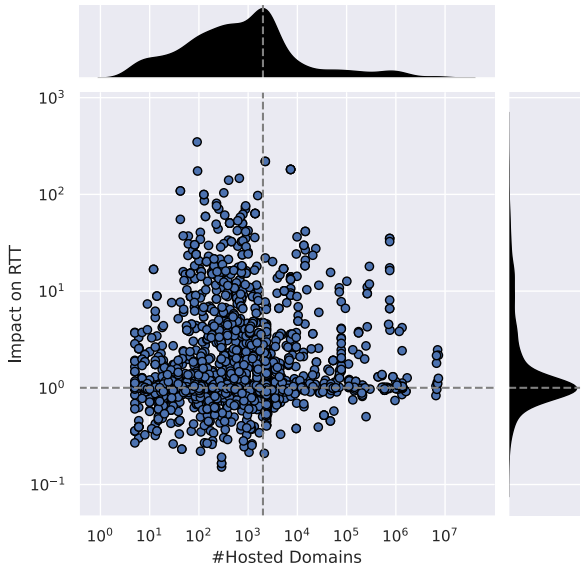


**Figure 8: RTT Impact vs. number of hosted domains. Most attacks had negligible impact on DNS resolution performance. However, some attacks led to peaks of 100-fold increase in resolution, likely disrupting end users.**

popular domain (*i.e.,* queried frequently, available in most caches) with a high TTL value may be less affected than a less popular one.

We also consider the targeted port for successful attacks and ask: *are successful attacks more likely to specifically target DNS service ports?* Recall from Section 6.2 that the majority of attacks do not target port 53. When we look at the port distribution of successful attacks, we see that the port distribution looks different: 49% of attacks target port 53 (DNS), 31% target port 80 (HTTP) and 11% target port 443 (HTTPS). While the fraction of attacks that are successful (i.e., they lead to resolution failures) is small, the difference in port distribution suggests that successful attacks are more likely to be specifically targeting the DNS. We speculate that this result is related to application-aware attacks, where attackers try to overload both the network and the application (*i.e.,* DNS authoritative software). Nevertheless, there are also other types of attacks that lead to a breakdown in name resolution, and there may be parallel attacks going on that we cannot observe through the network telescope.

*6.3.2 Resolution performance impairments.* Figure 8 shows the consequences of DDoS attacks in terms of RTT increase on different hosting sizes of NSsets. Most attacks did not cause observable impairments, but ≈ 5% of them (585) induced an 10-fold increase in RTT on 616 NSSets. In one-third (198) of these 585 attacks mentioned before, we see RTT peaking at more than 100-fold the baseline RTT. These high-impact attacks concentrated mainly on small-medium size infrastructure, hosting between 100-10K domain names. We also saw evidence of attacks on very large infrastructure (10M domain names); these manifested a smaller increase of 2-3 times the original RTT.

| Impact on RTT | Company |
|---|---|
| 348× | NForce B.V. |
| 219× | Co-Co NL |
| 181× | NMU Group |
| 174× | Hetzner |
| 146× | My Lock De |
| 140× | DigiHosting NL |
| 100× | Apple Russia |
| 76× | GoDaddy |
| 75× | Linode |
| 74× | ITandTEL |

**Table 6: The most affected companies in terms of RTT increase. The vast majority are small to medium size DNS hosting providers.**

Table 6 shows the most affected companies we investigated, by ASN and associated NS name. NForce B.V, a Dutch hosting provider was the most affected, followed by another Dutch company Co-Co NL. The third one Nordisk Media Utveckling is a Swedish company responsible for registration of popular and trademark-protected domains. We also found some general VPS/cloud providers (*e.g.,* Hetzner, My Lock, DigiHosting, Linode, ITandTEL) and large DNS hosting provider GoDaddy. Interestingly, we also found an attack against Apple Russia's ASN on Jan 21, 2022, well before Russia's February invasion of Ukraine and related attacks.

**Key Takeaway:** *Most attacks were ineffective, but some attacks had a critical impact, causing complete failure or dramatically increased latency of resolution.*

## 6.4 Attack Inferred Intensity Correlation

Correlating attack intensity inferred by the network telescope with impact on DNS infrastructure is non-trivial. In some cases the network telescope may observe only one low-intensity vector of a high-impact multi-vector attack. Or vice-versa, high-intensity attacks targeting large and redundant infrastructures may have little impact. Although the case studies we examined were clearly observable in the network telescope, overall we did not see using Pearson coefficient a strong correlation between RSDoS impact metrics and observable impacts on DNS resolution performance (Figure 9). We

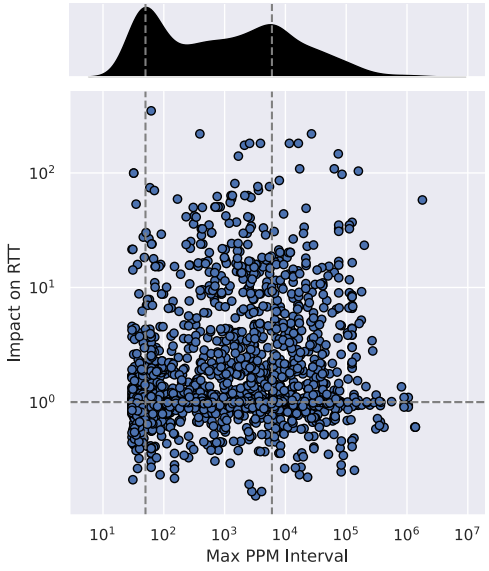


**Figure 9: We found low correlation between RTT Increase and Attack Intensity, implying that infrastructure handling capacity and deployment of resilience techniques play a fundamental role in withstanding DDoS attacks. Telescope data serves as a useful signal of ongoing attacks, and as an indicator of where to perform additional measurements, without providing exhaustive information on attack intensity.**

also saw high bandwidth (high packets/min) attacks on DNS infrastructure that continued to operate well. Attacks with low intensity as inferred by the telescope sometimes matched higher spikes in RTT of domain resolution. We speculate two underlying causes: deployment of resilience techniques, which mitigates the performance degradation induced by attacks; and multi-vector attacks not fully detected by the telescope. We also found no correlation between the RSDoS-reported number of attackers and DNS resolution performance impact or failure. We saw a bimodal distribution centered around 50 PPM (inferred to be 17K PPM after interpolating from the telescope address space to the entire IPv4 space) and 6000 PPM (inferred 2M PPM).

**Key Takeaway:** *Telescope data reveals signaling of ongoing attacks but does not enable prediction of performance impact.*

## 6.5 Attack Duration Correlation

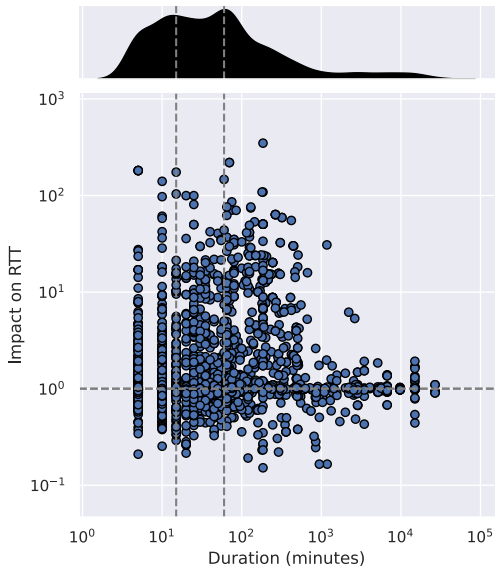


**Figure 10: Correlation between RTT Increase and Attack Duration. Attacks are generally short lived, but the longer the attack lasts, the more likely RTT increases will impair performance for end users.**

Figure 10 correlates attack impacts on resolution time and duration. The attack durations show a bi-modal distribution with modes at 15 minutes and 1 hour. High-impact attacks characterize these two intervals of the distribution, whereas attacks with longer duration have decreasing impact. Attacks may be short duration for several reasons, including that the attack succeeds and impedes responses that serve as backscatter signal, or that part of the attack is not visible to the network telescope. Long-term ineffective attacks could just represent background Internet noise. The only exception that we found was an attack against a German cloud provider, Contabo, which lasted 19 hours with a 30X increase in resolution RTT.

**Key Takeaway:** *Attacks with impacts on DNS are generally short-lived with an average duration between 15-60 minutes.*

## 6.6 Resilience technique efficacy

The impact of a DDoS attack is strongly related to the resilience techniques deployed. We analyzed several DNS resilience techniques to identify their possible effects on attack mitigation.

*6.6.1 Anycast vs DDoS.* Figure 11, Anycast deployments tend to suffer less under attack, (*i.e.,* RTT increase 1-1.5). Partial anycast deployments (*i.e.,* anycast deployed only on a subset of authoritative nameservers) show attacks having a small impact on the infrastructure. Most effective attacks are related to nameservers

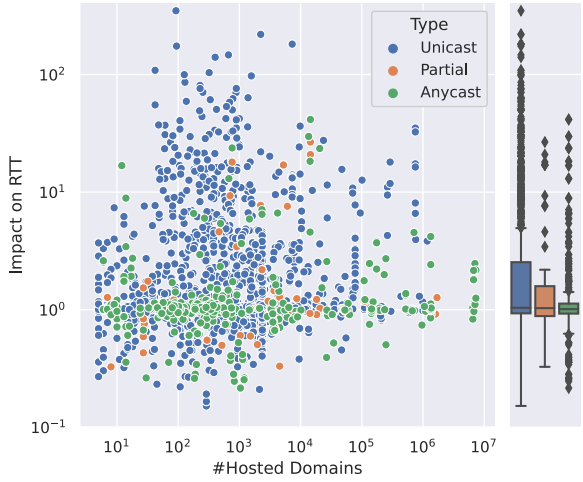


**Figure 11: Efficacy of anycast as a resilience technique. The impact of attacks on RTT increase for unicast hosted domains is generally higher. No DNS infrastructure experiencing 100-fold RTT increase was using anycast.**

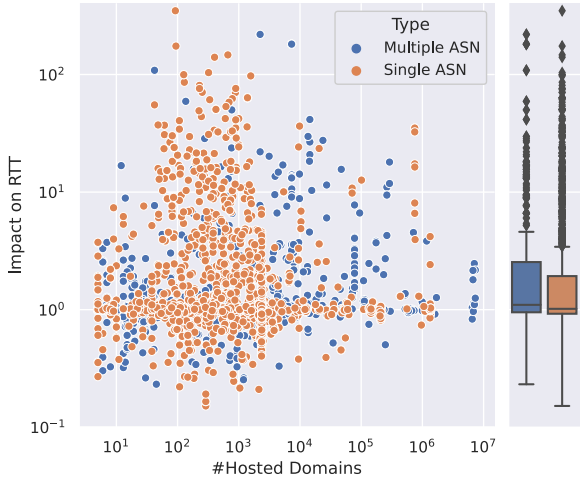


**Figure 12: AS Diversity efficacy as a resilience technique. NSsets that host a higher number of domains are more likely to have multi-AS deployments, but these alone do not provide a significant level of protection compared to single-AS deployments.**

hosted on unicast infrastructure (§6.3). In most cases of resolution failure, the domains relied on a unicast deployment. This result lends further support for the best practice of using anycast as a resilience technique against DDoS attacks.

*6.6.2 AS Diversity.* We did not find a clear link between AS diversity and effectiveness against DDoS attacks, but it seems to make more of a difference for larger ASNs (~1M domains), as shown in Figure 12. The graph shows the single behavior of a multi-variable system in terms of resilience technique (Anycast, AS Diversity, Prefix Diversity), which generally combine to reduce the impact of attacks. However, in cases of complete failure of reachability of domains, most of those domains (81%) relied on a single ASN deployment (§6.3).

*6.6.3 24 Prefix Diversity.* Nameservers hosted on a single /24 prefix are likely hosted on the same network infrastructures (*i.e.,* L2 switch, upstream router, etc). Figure 13 shows that using a single /24 prefix is generally the worst approach for deploying nameservers. Using two or more prefixes contributes significantly to resilience. §6.3 showed that 60% of NSsets that experienced failures were NSsets that relied on a single prefix. For most domains in this subset that experienced a complete failure in resolution, where 100% of queries failed to resolve, 30% of their NSsets were deployed on two prefixes and just 10% of their NSsets were served by three or more prefixes.

**Key Takeaway:** *Anycast deployments suffer less from attacks, indicating increased DNS infrastructure resilience. Hosting nameservers across multiple prefixes or multiple ASNs also appears to provide increased resilience to devastating attacks.*

## 7 RELATED WORK

*DDoS Attack Detection.* Several studies focused on inferring DDoS attacks on the Internet. Moore *et al.* introduced the method to detect randomly spoofed denial of service attacks (RSDoS) leveraged in this paper [22]. In this study, they explain that by monitoring a large address space, they can infer denial of service activity from the backscatter traffic observed in the Internet background radiation (IBR). Furthermore, they define threshold values to more accurately extract signals of attacks and eliminate sources of noise in the data. Leveraging this approach, Jonker *et al.* provided a macroscopic characterization of DDoS attacks on the Internet and investigated factors influencing migration to DDoS Protection Services (DPSs) [13]. Whereas Jonker *et al.* focused on hosting infrastructure, in this work we quantify and characterize the impact of attacks against DNS authoritative infrastructure.

Fachkha *et al.* designed an attack detection methodology using network telescopes by examining received DNS DDoS amplification attack traffic [10]. Other researchers have used honeypots to detect DDoS attacks.

Kramer *et al.* [17] developed AmpPot, a series of fake amplifier instances designed to monitor DDoS amplification and reflection attacks. Bailey *et al.* [4] designed an analogous system that used a two-tier approach with lightweight honeypots to monitor suspicious activity and high-end honeypots for behavioral analysis.

*Attacks against DNS infrastructure.* DNS infrastructure represents a frequent target of DDoS attacks. Moura *et al.* evaluated the

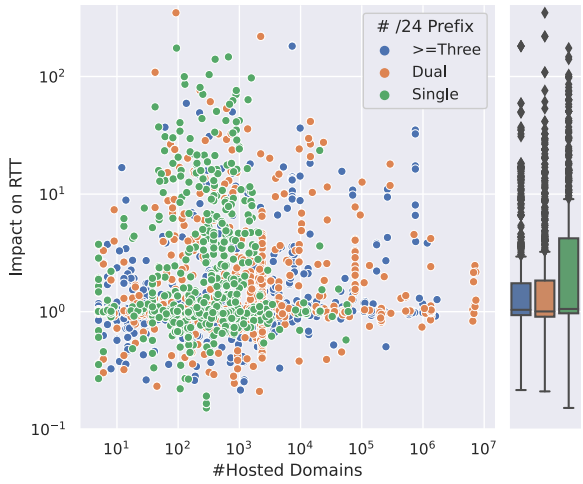


**Figure 13: /24 Prefix Diversity efficacy as a resilience technique. Our data set shows that using a single unicast prefix to serve DNS infrastructure was likely worst decision in terms of resilience. See Figure 11 for evidence these were mostly unicast prefixes.**

impact of large DDoS attacks against the DNS root server infrastructure in November 2015 [23]. They investigated how different services respond to stress and the performance impact of policies and mechanisms deployed to handle the attack. Their analysis demonstrated the efficacy of anycast as a resilience technique against DDoS attacks. Our work corroborates their findings, showing that deployment of anycast for DNS nameserver infrastructure remains the best protection against DDoS attacks.

In 2018, Moura *et al.* studied the benefit of DNS caching for DNS services severely impacted by DDoS attacks [25]. Their controlled experiments showed that the presence of caching allowed almost all end users to tolerate attacks causing up to 50% of packet loss on the DNS infrastructure.

Several studies focused on a high-impact attack against DNS provider Dyn in October 2016 [11, 19, 31], characterizing its impact, effects on the global internet ecosystem, and DNS customer behavior after the attack. Abhishta et al. further investigated Dyn DNS customer behavior before and after this attack, showing that customers that relied heavily on a single company for their authoritative nameservers switched to use other servers after successful attacks [1].

Four years after the Dyn attack, Kashaf *et al.* investigated third-party dependencies of modern web services. They showed that 89% of the Alexa Top-100K websites still critically depend on a third-party DNS, CDN or CA provider, despite that this exclusive dependency was why the Dyn attacks had such far-reaching effects on users [16].

*DNS resilience.* Focusing on DNS resilience, Allman *et al.* analyzed the structural DNS robustness of the DNS authoritative ecosystem over 9 years, showing adoption of different resilience techniques by DNS operators [3]. Sommese *et al.* expanded on this topic by providing an extensive characterization of the adoption of anycast in DNS authoritative infrastructure, showing a massive adoption for half of the domains they measured. This rapid adoption was correlated with a few large DNS providers beginning to offer DNS anycast services. Their analysis investigated the impact of anycast adoption on deployment of other resilience techniques [32]. Akiwate *et al.* characterized the prevalence of lame delegations on the DNS ecosystem and their negative impact on resilience and performance [2] Finally, Sommese *et al.* investigated the spread of parent-child inconsistency and its potentially negative impact on DNS resilience [34].

## 8 ETHICAL CONSIDERATIONS

One ethical concern in studying the impact of volumetric attacks on critical infrastructure is to avoid measurements that induce additional burden to infrastructure under attack. For this reason, analyses mostly relied on existing operational data collections that have zero to no impact on the attacked infrastructure or end users, *e.g.,* RSDoS (based on passive traffic observations), OpenINTEL (a lightweight probing architecture), Prefix2AS (leveraging RouteViews BGP collection infrastructure). We also developed a reactive measurement system to measure DNS infrastructure inferred to be under attack. To avoid causing harm by performing these measurements, we limited our query rate to 50 domains every 5 minutes for each IP under attack. Moreover, the system distributes these 50 queries evenly across the 5 minutes interval, sending approximately one query every 6 seconds, a negligible load on nameserver infrastructure.

Another ethical concern is public exposure of IP addresses target of successful attacks. One might argue that exposing these IP addresses might increase the chance of future attacks against them. To overcome this concern, we decided not to expose IP addresses in this work but did mention the associated companies. The only exception we made relates to already public information on the Internet (*e.g.,* newspapers, tech reports).

## 9 CONCLUSION

Calls for adoption of techniques to support resilience of DNS infrastructure began decades ago, starting with classical topological redundancy as described in RFC2182 [29], and more recently with anycast techniques [24]. Our results, including the lack of correlation between inferred attack intensity and actual impact on DNS users, provide evidence to support the relative effectiveness of such techniques at a macroscopic level. Well-provisioned and strategically designed DNS nameserver infrastructure can withstand severe attacks with negligible impairments to end users. On the other hand, even small attacks can pose risks to infrastructure that neglect to architect resilience into their critical DNS infrastructure. Our analysis corroborates the importance and refines the prioritization of several actionable recommendations for DNS infrastructure providers:

- Distributed anycast deployment is generally the most effective method to operationally mitigate the effects of DDoS attacks on end users (§6.6).
- Classical DNS resilience strategies, i.e., prefix and AS diversity, also contribute to resilience, although in our data set these techniques appear to provide less benefit to resilience relative to anycast.
- It is sensible for operators of small nameserver deployments to rely on multiple third-party large infrastructures to provide backup resilience.

With regard to the overarching goal of achieving a collective understanding of the DDoS landscape for network operators, policy makers, and researchers, we have demonstrated the importance of continuous monitoring of the global DNS infrastructure, including correlating macroscopic feeds of attack inferences with active measurements that capture evidence of impaired user experience.

This work also suggests a natural future direction. Using these macroscopic measurement data sources to trigger active measurements of critical infrastructure under attack can lead to additional insights into resilience and failure modes of different components. For example, measuring all nameservers for a given domain upon evidence of an attack will provide a more effective indication of effect on whether and how end users experience resolution failure. Measurement from multiple vantage points will also improve fidelity of inferences in the face of increasing anycast deployment. These techniques would overcome OpenINTEL's limitations of using the default rather than NS-exhaustive resolution process, and doing so from a single vantage point. The tradeoff is operational cost and measurement overhead. We have prototyped such a reactive measurement platform, and plan to use it to demonstrate the feasibility of real-time characterization of DDoS attacks against the global DNS infrastructure.

## ACKNOWLEDGMENTS

We thank our shepherd Mark Allman and the anonymous IMC reviewers for their insightful suggestions and feedback. This work is partially funded by the NWO-DHS MADDVIPR project (Grant Agreement 628.001.031/FA8750-19-2-0004) and the EU CONCORDIA project (Grant Agreement 830927). This work included funding awards from U.S. NSF (OAC-2131987, CNS-2120399, CNS-1705024, CNS-2202288). *The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF, DHS, or the U.S. government.*

## REFERENCES

[1] Abhishta Abhishta, Roland Van Rijswijk-Deij, and Lambert J.M. Nieuwenhuis. 2018. Measuring the impact of a successful DDoS attack on the customer behaviour of managed DNS service providers. *Computer Communication Review* 48, 5 (2018), 70–76. https://doi.org/10.1145/3310165.3310175

[2] Gautam Akiwate, Mattijs Jonker, Raffaele Sommese, Ian Foster, Geoffrey M. Voelker, Stefan Savage, and KC Claffy. 2020. Unresolved Issues: Prevalence, Persistence, and Perils of Lame Delegations. In *Proceedings of the ACM Internet Measurement Conference (IMC '20)*. 281–294. https://doi.org/10.1145/3419394.3423623

[3] Mark Allman. 2018. Comments on DNS Robustness. In *Proceedings of the Internet Measurement Conference 2018 (IMC '18)*. 84–90. https://doi.org/10.1145/3278532.3278541

[4] Michael Bailey, Evan Cooke, David Watson, Farnam Jahanian, and Niels Provos. 2004. A hybrid honeypot architecture for scalable network monitoring. *Univ. Michigan, Ann Arbor, MI, USA, Tech. Rep. CSE-TR-499-04* (2004).

[5] Brian Barret. 2022. DDoS Attempts Hit Russia as Ukraine Conflict Intensifies. (Feb 2022). https://www.wired.com/story/russia-ukraine-ddos-nft-nsa-security-news/

[6] Robert Beverly and Arthur Berger. 2015. Server Siblings: Identifying Shared IPv4/IPv6 Infrastructure Via Active Fingerprinting. In *International Conference on Passive and Active Measurement*. 149–161. https://doi.org/10.1007/978-3-319-15509-$8_12$

[7] CAIDA. 2017. Inferred AS to Organization Mapping Dataset. https://www.caida.org/data/as-organizations/. (2017).

[8] CAIDA. 2020. Routeviews Prefix to AS mappings Dataset for IPv4 and IPv6. (2020). http://www.caida.org/data/routing/routeviews-prefix2as.xml

[9] CAIDA. 2022. UCSD Network Telescope Daily Randomly and Uniformly Spoofed Denial-of-Service (RSDoS) Attack Metadata. (2022). https://www.caida.org/catalog/datasets/telescope-daily-rsdos/

[10] Claude Fachkha, Elias Bou-Harb, and Mourad Debbabi. 2014. Fingerprinting Internet DNS Amplification DDoS Activities. In *2014 6th International Conference on New Technologies, Mobility and Security (NTMS)*. 1–5. https://doi.org/10.1109/NTMS.2014.6814019

[11] Shane Greenstein. 2019. The Aftermath of the Dyn DDOS Attack. *IEEE Micro* 39, 4 (2019), 66–68. https://doi.org/10.1109/MM.2019.2919886

[12] Tijs Hofmans. 2020. Providers zijn maandag opnieuw getroffen door ddos-aanvallen [Dutch]. (Dec 2020). https://tweakers.net/nieuws/175228/providers-zijn-maandag-opnieuw-getroffen-door-ddos-aanvallen.html

[13] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. 2017. Millions of Targets under Attack: A Macroscopic Characterization of the DoS Ecosystem. In *Proceedings of the 2017 Internet Measurement Conference (IMC '17)*. 100–113. https://doi.org/10.1145/3131365.3131383

[14] Mattijs Jonker, Aiko Pras, Alberto Dainotti, and Anna Sperotto. 2018. A First Joint Look at DoS Attacks and BGP Blackholing in the Wild. In *Proceedings of the Internet Measurement Conference 2018 (IMC '18)*. 457–463. https://doi.org/10.1145/3278532.3278571

[15] Mattijs Jonker, Anna Sperotto, Roland van Rijswijk-Deij, Ramin Sadre, and Aiko Pras. 2016. Measuring the Adoption of DDoS Protection Services. In *Proceedings of the 2016 Internet Measurement Conference (IMC '16)*. 279–285. https://doi.org/10.1145/2987443.2987487

[16] Aqsa Kashaf, Vyas Sekar, and Yuvraj Agarwal. 2020. Analyzing Third Party Service Dependencies in Modern Web Services: Have We Learned from the Mirai-Dyn Incident?. In *Proceedings of the ACM Internet Measurement Conference (IMC '20)*. 634–647. https://doi.org/10.1145/3419394.3423664

[17] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. 2015. AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. In *Proceedings of the 18th International Symposium on Research in Attacks, Intrusions, and Defenses - Volume 9404 (RAID 2015)*. Springer-Verlag, Berlin, Heidelberg, 615–636. https://doi.org/10.1007/978-3-319-26362-$5_28$

[18] Jay Kreps. 2011. Kafka : a Distributed Messaging System for Log Processing.

[19] Yujing Liu, Zhilin Wang, and Nan Li. 2018. Characterizing the impact of DDoS attack on inter-domain routing system: a case study of the Dyn cyberattack. *Advances in Computer Science Research* 80, Csece (2018), 79–82. https://doi.org/10.2991/csece-18.2018.17

[20] Messaging, Malware and Mobile Anti-Abuse Working Group ( M3AAWG ). 2020. M3AAWG State of the Union: DDOS Attacks. (2020). https://www.m3aawg.org/blog/m3aawg-state-of-the-union-ddos-attacks

[21] Paul Mockapetris. 1987. Domain names - concepts and facilities. RFC 1034. (Nov. 1987). https://doi.org/10.17487/RFC1034

[22] David Moore, Colleen Shannon, Douglas J. Brown, Geoffrey M. Voelker, and Stefan Savage. 2006. Inferring Internet Denial-of-Service Activity. *ACM Trans. Comput. Syst.* 24, 2 (2006), 115–139. https://doi.org/10.1145/1132026.1132027

[23] G.C.M. Moura, R. de Oliveira Schmidt, J. Heidemann, W. B. de Vries, M. Muller, L. Wei, and C. Hesselman. 2016. Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In *Proceedings of the 2016 Internet Measurement Conference (IMC '16)*. 255–270. https://doi.org/10.1145/2987443.2987446

[24] Giovane Moura, Wes Hardaker, John Heidemann, and Marco Davids. 2022. Considerations for Large Authoritative DNS Server Operators. RFC 9199. (March 2022). https://doi.org/10.17487/RFC9199

[25] Giovane C. M. Moura, John Heidemann, Moritz Müller, Ricardo de O. Schmidt, and Marco Davids. 2018. When the Dike Breaks: Dissecting DNS Defenses During DDoS. In *Proceedings of the Internet Measurement Conference 2018 (IMC '18)*. 8–21. https://doi.org/10.1145/3278532.3278534

[26] NLNet Labs. 2020. Unbound DNS resolver. (2020). https://nlnetlabs.nl/projects/unbound/about/

[27] TransIP NOC. 2021. DDoS Post Mortem. (Mar 2021). https://tinyurl.com/transip-ddos

[28] Panda Security. 2020. Network Attacks: DoS and DDoS Attacks. (2020). https://www.pandasecurity.com/en/security-info/network-attacks/

[29] Michael A. Patton, Scott O. Bradner, Robert Elz, and Randy Bush. 1997. Selection and Operation of Secondary DNS Servers. RFC 2182. (July 1997). https://doi.org/10.17487/RFC2182

[30] Neil Rubenking. 2022. I Went to a Russian Website and All I Got Was This Lousy Teapot. (Feb 2022). https://www.pcmag.com/news/i-went-to-a-russian-website-and-all-i-got-was-this-lousy-teapot
[31] James Scott Sr and Winter Summit. 2016. Rise of the machines: The dyn attack was just a practice run december 2016. *Institute for Critical Infrastructure Technology, Washington, DC, USA* (2016).
[32] Raffaele Sommese, Gautam Akiwate, Mattijs Jonker, Giovane Moura, Marco Davids, Roland Martijn van Rijswijk - Deij, Geoffrey M. Voelker, Stefan Savage, Kimberley Claffy, and Anna Sperotto. 2021. Characterization of Anycast Adoption in the DNS Authoritative Infrastructure. In *5th Network Traffic Measurement and Analysis Conference, TMA 2021*. IFIP.
[33] Raffaele Sommese, Leandro Bertholdo, Gautam Akiwate, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, KC Claffy, and Anna Sperotto. 2020. MAnycast2: Using Anycast to Measure Anycast. In *Proceedings of the ACM Internet Measurement Conference (IMC '20)*. https://doi.org/10.1145/3419394.3423646
[34] Raffaele Sommese, Giovane C. M. Moura, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, K. C. Claffy, and Anna Sperotto. 2020. When Parents and Children Disagree: Diving into DNS Delegation Inconsistency. In *Passive and Active Measurement*. Springer International Publishing, Cham, 175–189. https://doi.org/10.1007/978-3-030-44081-$7_1$1
[35] Daan van Monsjou. 2021. Internetproviders zijn getroffen door ddos-aanvallen - update 2 [Dutch]. (Mar 2021). https://tweakers.net/nieuws/179580/internetproviders-zijn-getroffen-door-ddos-aanvallen.html
[36] Roland M. van Rijswijk, Mattijs Jonker, Anna Sperotto, and Aiko Pras. 2016. A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements. *IEEE journal on selected areas in communications* 34, 6 (June 2016), 1877–1888. https://doi.org/10.1109/JSAC.2016.2558918
[37] VirusTotal. 2022. VirusTotal. (2022). https://www.virustotal.com/
[38] Ramin Yazdani, Roland van Rijswijk-Deij, Mattijs Jonker, and Anna Sperotto. 2022. A Matter of Degree: Characterizing the Amplification Power of Open DNS Resolvers. In *Passive and Active Measurement: 23rd International Conference, PAM 2022, Virtual Event, March 28âĂŞ30, 2022, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 293–318. https://doi.org/10.1007/978-3-030-98785-$5_1$3
[39] Matei Zaharia, Reynold S. Xin, Patrick Wendell, Tathagata Das, Michael Armbrust, Ankur Dave, Xiangrui Meng, Josh Rosen, Shivaram Venkataraman, Michael J. Franklin, Ali Ghodsi, Joseph Gonzalez, Scott Shenker, and Ion Stoica. 2016. Apache Spark: A Unified Engine for Big Data Processing. *Commun. ACM* 59, 11 (oct 2016), 56–65. https://doi.org/10.1145/2934664