AFRL-RI-RS-TR-2022-059

MAPPING DNS DDOS VULNERABILITY TO IMPROVE PROTECTION AND PREVENTION

UNIVERSITY OF CALIFORNIA

MARCH 2022

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE

■ **AIR FORCE MATERIEL COMMAND** ■ **UNITED STATES AIR FORCE** ■ **ROME, NY 13441**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nations. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RI-RS-TR-2022-059 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE CHIEF ENGINEER:

/ S /
TODD N. CUSHMAN
Work Unit Manager

/ S /
JAMES S. PERRETTA
Deputy Chief
Information Warfare Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
		START DATE	END DATE
MARCH 2022	FINAL TECHNICAL REPORT	FEBRUARY 2019	FEBRUARY 2022

4. TITLE AND SUBTITLE

MAPPING DNS DDOS VULNERABILITY TO IMPROVE PROTECTION AND PREVENTION

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
FA8750-19-2-0004	N/A	62788F
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**
		R2P8

6. AUTHOR(S)

Kimberly Claffy, Alberto Dainotti, Mattijs Jonker, Roland Van Rijswijk-Deij, Raffaele Sommese, Anna Sperrotto, Elena Yulaeva

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of California San Diego
9500 Gilman Dr.
La Jolla CA 92093

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
Air Force Research Laboratory/RIGA 525 Brooks Road Rome NY 13441-4505	RI	AFRL-RI-RS-TR-2022-059

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The main goal of the project was to provide a measurement-based view of the DDoS-related threat landscape facing the Domain Name System (DNS) infrastructure, and to generate actionable intelligence enabling real-world improvements to the resilience of the DNS infrastructure against attacks. The project consisted of two pillars: (1) identifying DNS single points of failure and vulnerabilities and (2) mapping the DNS Distributed Denial of Service (DDoS) ecosystem. The accomplishments of the project include: (1) development of a methodology for detecting Anycast prefixes on the global Internet (Manycast2); (2) detailed analysis of Anycast deployment of DNS nameserver infrastructure, (3) development of DNSAttackStream, the software platform that enables a live view of the impact of spoofed DDoS attacks on the global DNS ecosystem by joining the CAIDA Network Telescope Reflected Spoofed Denial of Service (RSDOS) attacks data with live DNS measurement performed by OpenINTEL, an active DNS measurement project. The platform assists with identification of misconfigurations, vulnerabilities, and attacks, and (2) actionable recommendation for DNS operators. The intelligence and tools generated by the MADDVIPR project aid protection of the DNS and facilitate prevention of attacks against the DNS.

15. SUBJECT TERMS

DNS, DDoS, measurement infrastructure, cybersecurity, attacks

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	C. THIS PAGE		
U	U	U	SAR	24

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
TODD CUSHMAN	N/A

Contents

List of Figures

1.0 SUMMARY

The main goal of the project was to provide a measurement-based view of the DDoS-related threat landscape facing the Domain Name System (DNS) infrastructure, and to generate actionable intelligence enabling real-world improvements to the resilience of the DNS infrastructure against attacks.

The project consisted of two pillars: (1) *identifying DNS single points of failure and vulnerabilities* and (2) *mapping the DNS Distributed Denial of Service (DDoS) ecosystem*. These pillars provided two complementary views of the DNS DDoS problem. Execution of the first pillar enabled us to identify Single Points of Failure (SPoF) and vulnerabilities that can be exploited with DDoS attacks against nameserver infrastructure. From Pillar 2, we gained an overview of what is attacked in practice, based on continuous network measurements. These complementary views allowed us to synthesize a unified view of the DNS DDoS ecosystem, SPoFs and vulnerabilities and to create actionable intelligence for DNS protection and attack prevention.

The accomplishments of the project include: (1) development of a methodology for detecting Anycast prefixes on the global Internet (Manycast2); (2) detailed analysis of Anycast deployment of DNS nameserver infrastructure, (3) development of *DNSAttackStream*, the software platform that enables a live view of the impact of spoofed DDoS attacks on the global DNS ecosystem by joining the CAIDA Network Telescope Reflected Spoofed Denial of Service (RSDOS)[1] attacks data with live DNS measurement performed by OpenINTEL[2], an active DNS measurement project.

The main project deliverables are (1) the MADDVIPR framework consisting of a reactive DNS-measurement platform and a variety of dashboards that provide operators with insight into the current state of the DNS ecosystem. The platform assists with identification of misconfigurations, vulnerabilities, and attacks, and (2) actionable recommendation for DNS operators. The intelligence and tools generated by the MADDVIPR project aid protection of the DNS and facilitate prevention of attacks against the DNS.

[1] https://www.caida.org/catalog/datasets/telescope-daily-rsdos/

[2] https://openintel.nl/

2.0 INTRODUCTION

The DNS forms part of the core of the Internet. DNS provides the vital function of translating human-readable domain names into Internet Protocol (IP) addresses, thus acting as the phone book of the Internet. It also serves as a support infrastructure for most applications, commercial content distribution platforms, and many security services [1]. DDoS attacks against the DNS can, therefore, have devastating effects – they are one of the most critical cyber-threats on the modern-day Internet. They are cheap, effective, and they keep growing in intensity as Internet connectivity and device capability grows. Securing the DNS against DDoS attacks is not a trivial task. Commercial solutions are expensive and can introduce a single point of failure by aggregating traffic toward a single entity.

To address these challenges, we proposed to step back and analyze what DNS infrastructure needs protection from what. First, we identified the DNS single points of failure and vulnerabilities suggesting how DNS nameservers can become the target of DDoS attacks. Then we provided a comprehensive overview of current DDoS attacks against the DNS by studying the attackers, attacks, and targets. Finally, we prototyped the MADDVIPR framework that provides a coherent, unified view of the DNS DDoS ecosystem. Taken together, our efforts have yielded actionable information on how to improve DNS resilience against DDoS attacks.

3.0 METHODS ASSUMPTIONS, AND PROCEDURES

Our approach comprised three tasks: (i) Identifying DNS single points of failure and vulnerabilities, which aims at identifying how the DNS can become the target of future DDoS attacks; (ii) Analyzing the DNS DDoS ecosystem, which provides a measurement-based overview of current DDoS attacks against the DNS by studying the attackers, the attacks and targets; and (iii) developing the MADDVIPR framework -- a coherent, unified view of the DNS DDoS ecosystem and the DNS single points of failure and vulnerabilities, that yields actionable information for operators on how to improve their own DNS resilience against DDoS attacks.

A core part of our approach is that we used the unique datasets that project partners have created, maintain, and use. The University of Twente contributed data from the OpenINTEL project, which collects daily active measurements of 60% of the global DNS name space, including the main top-level domains such as .com, .net and .org, and a growing number of country-code Top-Level Domains (TLD), including the .nl ccTLD. OpenINTEL has been collecting data since February 2015, and thus provides a unique longitudinal view of the evolution of large parts of the DNS. CAIDA contributed data from the UCSD Network Telescope, a large (originally /8, now a /9 and /10) globally routed block of IPv4 address space that CAIDA monitors for incoming traffic. This instrumentation provides visibility into, among other things, backscatter from ongoing DDoS attacks.

3.1 Pillar I – Identifying DNS Single Points of Failure and Vulnerabilities

We considered two classes of DNS vulnerabilities: (i) single points of failure and (ii) vulnerabilities due to misconfigurations or suboptimal configurations. We defined a *single point of failure* as the situation where authoritative information for a domain is available from only a single DNS operator, i.e., there is no redundant source in case this single source is unreachable. Potential *misconfigurations*, or suboptimal configurations are, for example, mismatches between DNS delegations in the parent and child zone or cross-domain vulnerabilities (if one domain is attacked, others are also affected as collateral damage), etc.

Our approach to identifying these vulnerabilities consisted of two steps:

(1) Identifying single points of failure. Using data collected by the OpenINTEL project and other topological information (e.g. the autonomous systems or IP prefixes that host name serves) we were able to identify single points of failure by mapping authoritative name servers back to their operators. The main challenge we overcame was to compose a set of views of OpenINTEL data that combined these different ways to identify operators, but still yielded a consistent and coherent view of single points of failure.

(2) Identifying misconfigurations and suboptimal configurations. We first performed a systematic analysis of both good practices in terms of configuring DNS for domains, and of common configuration errors. Based on this analysis, we used longitudinal data collected by the OpenINTEL project to quantify the occurrence of misconfigurations, and to analyze whether we can observe trends in the frequency at which these misconfigurations occur. The main challenge here was to define suitable signatures of such misconfigurations to detect in the sizable datasets from OpenINTEL.

We went one step further and considered the potential impact of our approaches to minimize or mitigate DNS vulnerabilities. Once we identified common vulnerabilities, we surveyed the

state-of-the-art in terms of standardized or proposed approaches to improve DNS resilience and analyzed to what extent these approaches address the vulnerabilities we discovered.

3.2 Pillar II – Mapping the DNS DDoS Ecosystem

A thorough understanding of the characteristics of DDoS attacks on the DNS plays an essential role in both attack prevention and effective protection. For this reason, in this project we devised a methodology that maps the DNS DDoS attack ecosystem. The mapping involves a macroscopic analysis of data on past and present attacks. Initially we based our method on two data sources:

(1) The UCSD Network Telescope that offers an excellent vantage point to capture trace data of DDoS attacks in which attackers try to disguise the source of malicious network traffic by applying (uniformly) random IP spoofing [2]. CAIDA has analyzed the observable "backscatter" traffic reaching the Telescope as a result of such (D)DoS attacks for many years, allowing us to assess historical trends.

(2) Other previously proven data sources to account for attack types that do not involve uniformly random spoofing. One example is data from the AmpPot project [3], which leverages honeypots to capture traces of (D)DoS attacks that involve the abuse of reflectors.

Where opportunities arose, we augmented our analysis with other sources of trace data such as evidence from botnet commands and control servers. We integrated the resulting methodology into a new software platform *DNSAttackStream* [4], the software that automatically analyzes the DNS (D)DoS attack ecosystem. *DNSAttackStream* generates (near) real-time intelligence on ongoing attacks as they show up in the Network Telescope and other data sources, and identifies attack sources, targets, and characteristics.

3.3 Synthesizing a Unified View of the DNS DDoS Ecosystem

Pillars (1) and (2) provide two complementary views of the DNS DDoS problem. From one side (Pillar (1)) we are now able to identify SPoF and vulnerabilities that can be exploited in case of DDoS attack against the DNS; from another side (Pillar (2)), we have gained an overview of what is attacked in practice, based on continuous network measurements. The next step was to synthesize a unified view of the DNS DDoS ecosystem, SPoFs and vulnerabilities in order to create actionable intelligence for DNS protection and attack prevention. This stage of the project concentrated on the following activities:

(1) Identification of the impact of possible attacks. By combining knowledge of attack targets, attack trends and SPoF, we are now able to determine the impact an attack could potentially have on the DNS infrastructure and collateral damage on other services.
(2) A view of future attacks. By combining the DNS DDoS ecosystem with knowledge of the vulnerabilities identified in Pillar (1), we can identify weak points in the practical use of the DNS and its configuration that might lead to future attacks.
(3) Prioritization of risks. We explored how this combined knowledge can be used to create a clear prioritization and ranking of SPoF and vulnerabilities that are a major risk for the DNS, and guide operators and security experts in attack mitigation and prevention.

4.0 RESULTS AND DISCUSSIONS

4.1 Identifying DNS Single Point of Failure and Vulnerabilities

Background Literature Survey of Methodologies

First, we focused on understanding vulnerabilities and misconfigurations in the DNS ecosystem. We performed an overview study of the relevant literature to identify existing contributions. We reviewed the current state of the art in DNS vulnerabilities and attacks, including reflection and amplification attacks. One use of the DNS as an attack vector is to spread amplification and reflection attacks. DDoS attacks create network congestion on paths to targets of the attacks. In combination with spoofing, DDoS attacks abuse User Datagram Protocol (UDP)-based services to cause reflection and trigger large responses from the UDP services thus achieving amplification. DDoS Reflection and Amplifications Attacks are often carried out using open resolvers. We systematically described all classes of vectors of DNS amplification and provided a brief description of CAIDA's Spoofer project [5] as an example of spoofing attacks mitigation. We provided an overview of current exploitable attacks against the DNS including DNS cache poisoning, domain hijacking, random subdomain attacks, NXDOMAIN attacks, phantom domain attacks, NSED White Lie DoS, DDoS DNS flood attacks, and Distributed Reflection Denial of Service. We described some frequent misconfigurations and vulnerabilities present in the DNS and analyzed in the literature. Those included misconfiguration of DNS Security Extension (DNSSEC), parent-child zone mismatch, single point of failure (e.g. single or duplicated nameserver (NS) records, infrastructural single point of failure), and dangling pointer misconfiguration. We also have briefly discussed several current measures to reduce the attack exposure of the DNS, and their adoption. We submitted our technical report [6] to DHS and published it online.

This background analysis was fundamental for our research since it revealed methodologies for how to measure the spread and the impact of these vulnerabilities and with possible countermeasures aimed at improving the resilience of DNS.

Investigation of identified vulnerabilities: Parent-child inconsistency

We investigated two of the identified vulnerabilities and misconfigurations in the DNS system: Parent-Child Inconsistency in the DNS hierarchy and Orphan and Abandoned Records.

In the first case, we studied the consistency of the replicated information along the DNS hierarchy between parent and child zone. DNS is a hierarchical, decentralized, and distributed database. A key mechanism that enables the DNS to be hierarchical and distributed is delegation of responsibility from parent to child zone*s*—typically managed by different entities. According to RFC1034 [7], authoritative NS records at both parent and child should be "consistent and remain so", but we find inconsistencies for over 13M second-level domains. We classified the type of inconsistencies we observed, and the behavior of resolvers in the face of such inconsistencies, using RIPE Atlas[3] to probe our experimental domain configured for different

[3] https://atlas.ripe.net/

scenarios. Our results [8] underlined the risk such inconsistencies pose to the availability of misconfigured domains. We presented our research and recommendations at Passive and Active Measurements Conference 2020 and at RIPE-80.

We Developed a SuperDNS tool for detection of DNS parent-children misconfiguration in a controlled environment [9] which provides insight to the operators and users.

Investigation of identified vulnerabilities: Orphan records

We quantified the *orphan records* misconfiguration, in which a glue record for a delegation that does not exist anymore is forgotten in the zone file. *Orphan records* are a security hazard to third-party domains that have these records in their delegation, as an attacker may easily hijack such domains by registering the domain associated with the orphan. We extended the previous work by Kalafut et al [10] by identifying a new type of glue record misconfiguration - which we refer to as *abandoned records* - and by performing a broader characterization. We discovered that for the .com and .net TLDs, the number of orphan records has fallen to zero, which means that operators have introduced mechanisms for cleaning their zone files. Unfortunately, not all TLD registry operators have adopted these best practices. For some TLDs, the number of orphan records have increased over 10 years. Also, in the new generic Top-Level Domains (gTLDs) this misconfiguration is widespread.

We also discovered and analyzed another misconfiguration, the abandoned record. Our analysis showed that this misconfiguration is broader than the orphan one. Common sense would registries or registrars should remove abandoned records, as they potentially represent the initial stage of orphan creation. Our study also showed that the removal of these records from the zone file may not be a simple operation since it can incur the risk of breaking other domains. We recommended that registries should address the nature of the resources related to orphan records (i.e., hosted websites or domains) and of their related traffic by actively registering these domains and intercepting them. Finally, we suggested that all registry operators address this misconfiguration by at least making domains related to orphan records not available for registration or by considering cleaning up their zone removing orphans. We published and presented our findings and recommendations at 2*020 IEEE European Symposium on Security and Privacy Workshops*, OARC33 and WTMC2020 [11].

Investigation of identified vulnerabilities: Lame delegations

With UCSD collaborators we performed a comprehensive measurement study of *lame delegations*, using both longitudinal zone data and active querying. *Lame delegations* occur when a nameserver responsible for a domain is unable to provide authoritative information about it. They introduce performance and security risks. Using comprehensive collections of active and passive DNS measurements (covering 49 M and 499 M domains respectively), we found that lame delegations are surprisingly common: roughly 14% of registered domains that we actively measured had at least one lame delegation, and most of those had no working authoritative nameservers. Even for domains with working alternative nameservers, our measurements show that these lame delegations impair DNS performance (average resolution latency increased by 3.7×) in addition to producing substantial unnecessary load on existing nameservers. Finally, we found that unregistered or expired domains in lame delegations can create significant security

risk. We identified at least three instances over the last nine years in which an attacker could have hijacked thousands of domains by registering a single nameserver domain.

Analysis of this phenomenon led us to discover an unforeseen interaction between registrar practice and the constraints of registry provisioning systems that has inadvertently made hundreds of thousands of domains vulnerable to hijacking due to accidental lame delegations. This practice has persisted for over twenty years. We worked with registries and registrars to try to remediate this vulnerability for existing domains, and discussed ways to remediate it in the longer term (which requires changes to registrar business practices or protocols that seem unlikely without some incentive to do so).. We are exploring ways to combine daily zone data and periodic active measurements to automatically identify and report lame delegations as they are created. We published and presented these findings at IMC2020 [12].

4.2 Analyzing DNS Vulnerabilities and the DNS DDoS Ecosystem

Identifying DNS resilience vulnerabilities: Broad assessment of anycast deployment

We investigated and identified DNS configuration mechanisms that are likely to improve or degrade the resilience of the DNS to DDoS attacks. Using OpenINTEL and new measurement experiments we designed and executed, we quantified the observable deployment of these mechanisms.

Anycast addressing (assigning the same IP address to multiple, distributed devices) has become a fundamental approach to improving the resilience and performance of Internet services, but its conventional deployment model makes it impossible to infer from the address itself that it is anycast. Existing methods to detect anycast IPv4 prefixes present accuracy challenges stemming from routing and latency dynamics, and efficiency and scalability challenges related to measurement load. We reviewed these challenges and introduced a new technique we call "*MAnycast2*" that can help overcome them. This technique uses a distributed measurement platform of anycast vantage points as sources to probe potential anycast destinations. This measurement methodology eliminates any sensitivity to latency dynamics, and greatly improves efficiency and scalability. We researched alternatives to overcome remaining methodological challenges relating to routing dynamics, suggesting a path toward establishing the capability to complete, in under 3 hours, a full census of which IPv4 prefixes in the ISI hitlist are anycast. We published and presented our methodology at IMC2020 [13].

We built on this method to create another important outcome of this project: the *Anycast Census dataset* [14]. This census dataset is derived by deploying MAnycast^2 on the SIDN Anycast network of 20 geographically distributed nodes. To provide a method of cross-validation, we integrated *iGreedy*[4], a tool that detects, enumerates and geolocates anycasts instances. We performed *iGreedy* measurements with a set of 500 RIPE Atlas probes equally geographically distributed (200 km minimum distances between them). SIDN and U.Twente currently update the Anycast Census dataset quarterly and share it in JSON format.

We used these anycast census datasets to perform our DNS DDoS resilience analysis, by quantifying the adoption of anycast to support authoritative domain name service for TLDs and SLDs. Comparing two comprehensive anycast census datasets in 2017 and 2021, with DNS

[4] https://www.ict-mplane.eu/public/igreedy-anycast-enumeration-and-geolocation-module

measurements captured over the same period, we found a high rate of adoption of anycast as a resilience mechanism, reaching 97% for TLDs and 62% for SLDs. This adoption is driven mostly by engineering choices of a few very large DNS infrastructure providers. In our dataset, one provider (GoDaddy) was responsible for the majority of anycast adoption in SLDs. We also examined the relationship of anycast deployments to other traditional metrics of infrastructure diversity.

Our findings show that anycast adoption changes the DNS service availability risk profile but does not eliminate all resilience risks. In fact, anycast can hide certain types of availability failures, and limit recovery options. A mixed deployment that includes traditional unicast redundancy as well as anycast options mitigates this risk but increases cost and complexity. We discussed these aspects, and how the pervasive use of anycast merits a re-evaluation of how to measure DNS resilience. We published and presented this research at the Network Traffic Measurement and Analysis Conference (TMA, 2021) [15].

Identifying DNS resilience vulnerabilities: Broad assessment of anycast deployment

To understand how DDoS attacks affect DNS infrastructure, we developed the *DNSAttackStream* prototype [4], which provides a live overview of the impact of DDoS attacks on DNS infrastructure.

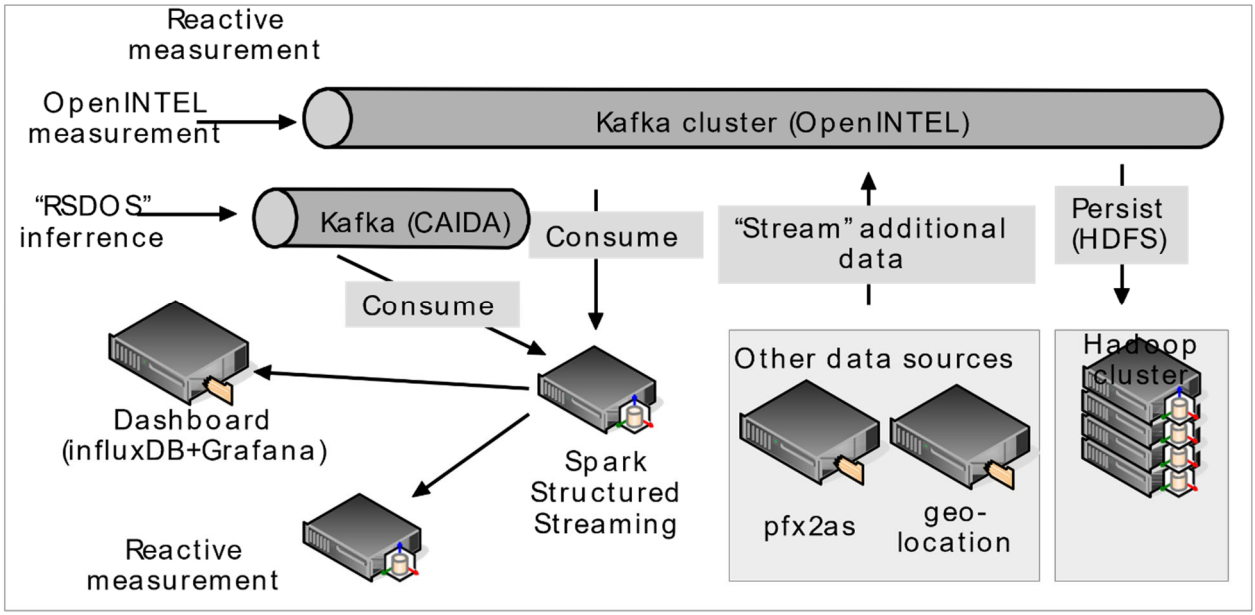

Figure 1 Data flow architecture of DNSAttackStream platform

DNSAttackStream represents the first step in integrating different data sources into the MADDVIPR framework. We started with Randomly Spoofed Denial-of-Service (RS-DOS) attack information, collected by the UCSD Network Telescope[5] project, and joined it with OpenINTEL live measurements. *DNSAttackStream* merges the list of IP addresses inferred to be under attack based on the UCSD Network Telescope data every 5 minutes with the list of IP addresses of authoritative nameservers measured by OpenINTEL. This mechanism allows us to provide insights into the number of authoritative nameservers and related Second Level Domains (SLDs) affected by attacks. To integrate these two sources of information, we implemented a streaming pipeline, using Kafka as a message broker for retrieving live data, a Spark streaming application for joining the two live datasets, Telegraf as middleware for InfluxDB, InfluxDB for storing time series of aggregated attack information, and finally Grafana for the implementation of a live dashboard. We documented the *DNSAttackStream* in the technical report [4] that we submitted to DHS.

4.3 MADDVIPR Framework Prototype

The core element of the MADDVIPR project is the MADDVIPR framework. Through this framework, we consolidated and synthesized our research on the DNS single points of failure, vulnerabilities, and DDoS attacks against the DNS to produce actionable intelligence for DNS operators (detect, analyze and prevent DDoS attacks), and to improve infrastructure resilience.

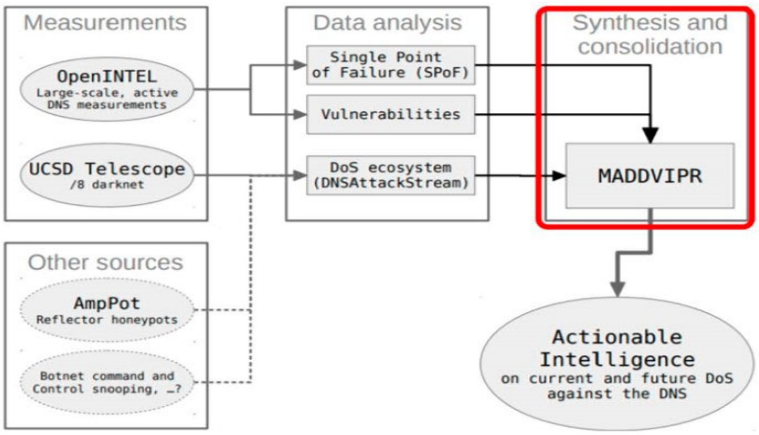

Figure 2. MADDVIPR Architecture

The MADDVIPR framework [16] consists of two main components: (1) a highly configurable, reactive DNS measurement platform, designed to be scalable through a cloud-based, multi-tenant infrastructure, geographically distributed over the global Internet and (2) a variety of dashboards, fed with real-time data and intelligence, to provide operators insights into the current state of the DNS ecosystem and to help them to identify misconfigurations, vulnerabilities, and attacks.

[5] https://stardust.caida.org/

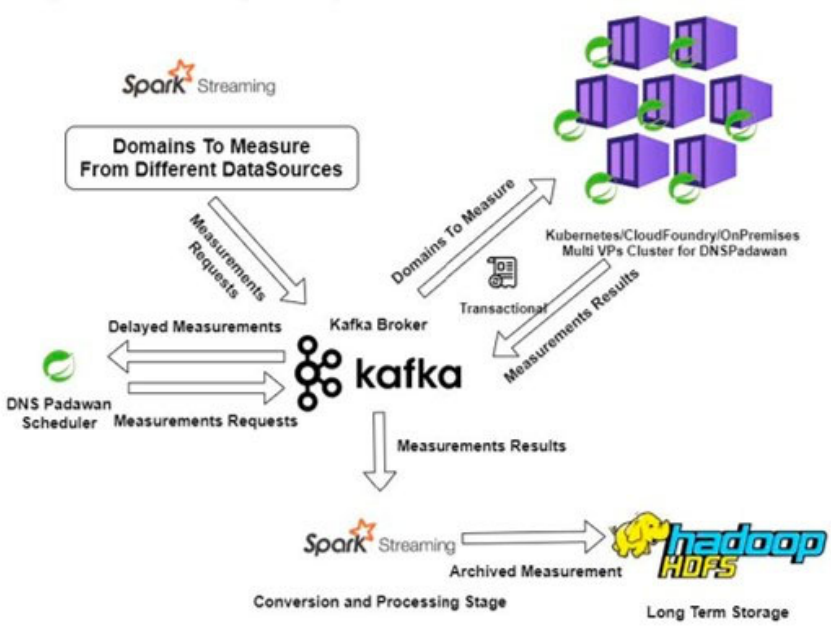

Figure 3. MADDVIRP measurement framework overview

Figure 3 illustrates the overall architecture of the measurement framework. The domain names to measure, which are inferred from different data sources, are published live on a Kafka Topic by a streaming application (e.g., in Spark Structured Streaming). Then, we differentiate the measurements into two main types: *Immediate* and *Scheduled*. The measurement software (the working name of which is *DNSPadawan*) performs *immediate* measurements in a best-effort approach as soon as possible. The software can be deployed on several platforms (Kubernetes, CloudFoundry, and on-premise's deployments). *Scheduled* measurements are collected by the *DNSPadawan* scheduler component, which is responsible for scheduling a measurement at the requested time and for the requested number of repetitions. By using the proven Kafka Streams technology we can gain out-of-the-box embedded load balancing mechanisms across different VPs. The *DNSPadawan* measurement and scheduler components are implemented using the Java Spring framework. This approach allows us to obtain a solid system, natively interoperable with Kafka and several cloud orchestration technologies.

4.4 Testing and Evaluation

We analyzed and processed large-scale datasets (OpenINTEL, Telescope, etc) using UCSD and U. Twente high-performance computing environments. We used our extensive contacts in the DNS operator community to solicit feedback on the results of our various studies and data collection platform. We leveraged the DNS Operations Analysis and Research Center (See Section 4.5) to get early feedback from organizations that benefit from MADDVIPR's results.

As an applied scientific research project, we used well-attended technical forums as a channel for knowledge transfer. We also published our studies at the premier scientific conference on Internet measurements, which leveraged peer review as a mechanism for evaluation of our methods, algorithms, and results. (see our publications and evaluation of a scalable method for Appendix A: Presentations and Publications).

4.5 Collaboration with Industry

Collaboration with industry (DNS and network operators) is our primary channel for validating our methods and inferences regarding misconfigurations and vulnerabilities in the DNS ecosystem. Communication with practitioners also helps us understand resilience mechanisms and technical decisions that embody tradeoffs among resilience, performance, and cost. An explicit goal of this project is to convey knowledge we obtain back to operators, to inform their strategies, policies, and operations to improve DNS resilience.

We used our contacts in the DNS operator community to solicit feedback on the results that emerge from this project. We got substantial feedback from the DNS Operations Analysis and Research Center (DNS-OARC). CAIDA is one of the founding members of DNS-OARC and the University of Twente is an academic member. DNS-OARC organizes two annual meetings where members presented work to and solicit feedback from fellow members. This membership includes all the big names in DNS operations, ranging from registry operators for top-level domains to large Internet brands.

We invested significant effort into informing the Internet community of the actionable recommendations of this project. Our work on parent-child misconfigurations led to a publicly available software tool [9], and suggestions on the Internet Engineering Task Force (IETF[6]) DNS Operations Working Group Internet-Draft aiming to solve the misconfiguration problem.

We reached out directly to DNS operators with actionable knowledge to address vulnerabilities arising from misconfigurations we discovered. Following our study, Afilias[7] informed registrars and registry clients that it would take steps to remove orphan glue records from 200+ TLD zones in its care. They specifically acknowledged our help with fixing misconfigurations in their zone files in their CircleID blog [17]. [8]

We opened the *DNSAttackStream* dashboard[9] to operators and demonstrated its ability to identify an attack to a large Dutch DNS provider and showing the impact of the attack on the DNS at the internal event for the Netherlands National Cyber Security Centre.

We also used Internet community social media channels to share the project's actionable outputs. For example, we used APNIC (Regional Internet Registry administering IP addresses for the Asia Pacific) blog to share the results of our research and evaluation of a scalable method for identification of anycast prefixes [18].

We have gained valuable experience in sharing data, information, and best practices with operators. We will continue to present our work at scientific and technical conferences, and transfer actionable knowledge and measurement technology to operators, notifying them of potential resilience problems and misconfigurations.

---

[6] The IETF is an international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. https://www.ietf.org/

[7] https://afilias.info

[8] CircleID is a platform for publishing Internet related articles and announcements. https://www.circleid.com

[9] http://192.87.172.248:3000/d/AOK3LzVMk/dnsattackstream?orgId=1

5.0 Conclusions and Actionable Recommendations

We performed a comprehensive analysis of the DDoS ecosystem targeting the DNS – attack sources, targets, and characteristics observed in DDoS attack traffic data, and assessed vulnerabilities and single points of failure that threaten the resilience of the DNS under such DDoS attacks. Combining these two perspectives we documented a clear view of the threat landscape facing DNS infrastructure, and generated recommendations enabling real-world improvements to the resilience of the DNS against attacks. We developed and started to operate a measurement and analysis system that generates actionable intelligence that can protect the DNS against DDoS attacks and provides recommendations for preventative measures that minimize the risk and impact of such attacks. Based on our research and collaboration with industry, we compiled a set of recommendation for operators. Below is a brief overview of our recommendations.

5.1 Keep DNS Parent and Children Zone Consistent

The first recommendation follows our study on parent-children delegation inconsistency [8]. In this study, we characterized the spread and the impact of the inconsistency of delegation (NS records) between TLDs and SLDs zones. We showed how this problem can lead to risks in terms of resilience, unavailability, privacy leaks, and hijacking. Therefore, our first recommendation for operators is: *Keep Parent and Children Zone redundant records consistent both in terms of set of all Resource Records (RRset) and Time to Live (TTL) values*.

5.2 Clean DNS Orphan Glue Records from Zone Files

The second recommendation follows our study on orphan and abandoned records [11]. In this study, we characterized the problem of orphan and abandoned glue records, showing how they represent a problem in terms of zone file pollution and how they can lead to the risk of hijacking events. The problem is extremely relevant for TLD operators. Therefore, our second recommendation for (TLD) operators is: *Keep zone files clean by removing all the unnecessary or expired glue records.*

5.3 Do Not Use Anycast as the Sole Mechanism for DNS Resilience

The third recommendation follows our study on anycast deployment in DNS authoritative infrastructure [15]. In the study, we mapped the adoption of anycast in the DNS infrastructure of more than 210 Million SLDs between 2017 and 2021, showing that half of them used anycast for part of their authoritative infrastructure. We found this adoption was mainly driven by large providers, and that in several cases, Anycast was used as the sole mechanism for DNS resilience. We characterized the adoption of other resilience mechanisms (AS, Prefix, IP, and GeoLocation diversity) showing that the most resilient scenarios leveraged both unicast and anycast deployments. Therefore, our third recommendation for DNS operators is: *Use and widely adopt anycast as a resilience mechanism for DNS in combination and not in substitution of other resilience mechanisms (AS, Prefix, IP, and Geolocation Diversity).*

5.4 Always Check Your Anycast Routing configuration

The fourth and final recommendation follows our study on anycast census [13]. In the study, we deployed a methodology for detecting anycast prefixes at scale. While performing the census we detected several networks that exhibit different behavior based on their routing configuration, resulting in some cases with routing all the traffic to a single location. Therefore, our fourth recommendation for DNS operators is: *Always check your Anycast routing configuration and the catchment of the different sites of your anycast deployments*.

5.5 Future Work Based on Outcomes

Both the U.S. and Netherlands teams are continuing research using the data sets compiled and supported with this DHS award. In the short term, we will complete the paper documenting the *DNSAttackStream* methodology and system and submit it to the ACM/SIGCOMM Internet Measurement Conference in May 2022. The Dutch team is working on a paper describing the revisions made to their OpenIntel platform, in part supported by this award. Their funding continues to late 2022. UC San Diego will host Twente PhD candidate Raffaele Sommese from May-August 2022, a visiting internship that we had planned for 2019 but the pandemic postponed it. We will use this time to: (1) advance our methods for reactive measurement to fingerprint and detect attackers based on telescope observations; (2) join the reverse DNS data set that U Twente has been collecting for 2 years with our other data sources, to analyze IPv4 coverage, infer semantics of the namespace, and compare its maintenance and structure with forward DNS mappings

In the medium term, the Dutch team will undertake work under its recent 2M euro project award for scientific research to start up the "Responsible Internet"[10] which will include promotion of operational practices that help safeguard against DDoS and other attacks against or leveraging DNS infrastructure

In parallel, UC San Diego is exploring revenue opportunities to sustain the operation of the UCSD network telescope instrumentation to support this and other cybersecurity research. In the meantime, this instrumentation is supported by short-term NSF and DARPA funding.

[10] https://www.sidnlabs.nl/en/news-and-blogs/three-more-things-you-need-to-know-about-the-responsible-internet

6.0 REFERENCES

[1] Jonker M., Sperotto A., van Rijswijk-Deij, R., Sadre, R., and Pras, A., "Measuring the Adoption of DDoS Protection Services," *Proceedings of the 2016 ACM IMC*, Santa Monica, CA, US, 2016 Nov 14, pp. 279-285.

[2] Moore, D., Shannon, C., Brown, D.J., Voelker, GM., and Savage,S., "Inferring Internet Denial-of-service Activity," *ACM Transactions on Computer Systems*, **24(2)**, 2006, pp. 115–139.

[3] Kr¨amer, L., Krupp, J., Makita, D., Nishizoe, T., Koide, T., Yoshioka, K., and Rossow, C., "AmpPot: Monitoring and Defending Against Amplification DDoS Attacks," *Proceedings of the Research in Attacks, Intrusions, and Defenses - 18th International Symposium, RAID 2015,* Kyoto, JP, 2015 Nov 2-4.

[4] MADDVIPR, "DNSAttackStream Prototype," University of Twente and CAIDA, UCSD , September 2020, URL: *https://maddvipr.org/deadline/documents/DNSAttackStream.pdf* . Accessed 15 February 2022.

[5] CAIDA, "Spoofer project," URL: *https://www.caida.org/projects/spoofer.* Accessed 15 February 2022.

[6] Sommese , R., Sperotto A., van Rijswijk-Deij, R., Dainotti, A., Claffy K., " Background research on DNS-related DDoS vulnerabilities,*"* University of Twente and CAIDA, UCSD, January 2019, URL: *https://www.caida.org/funding/usnl-maddvipr/background_research_dns_related.pdf* . Accessed 15 February 2022.

[7] Mockapetris, P., "Domain names - concepts and facilities," RFC 1034, IETF, November 1987. URL: *http://tools.ietf.org/rfc/rfc1034.txt* Accessed 15 February 2022.

[8] Sommese, R., Moura, G., Jonker, M., Rijswijk-Deij, R. V., Dainotti, A., Claffy, K. C., & Sperotto, A (2020) "When Parents and Children Disagree: Diving into DNS Delegation Inconsistency," In Sperotto A., Dainotti A., and Stiller B., "Passive and Active Measurement," *Proceedings of Passive and Active Measurement Conference 2020,* online, March 2020, *Lecture Notes in Computer Science*, vol 12048, Springer, Cham. https://doi.org/10.1007/978-3-030-44081-7_11

[9] University of Twente, "SuperDNS tool", URL https://superdns.nl/ Accessed 15 February 2022.

[10] Kalafut, J., Gupta, M., Cole, C.A., Chen. L., and Myers, N.E., "An empirical study of orphan DNS servers in the Internet," *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement 2010*, Melbourne, AU, November 2010, pp. 308-314.

[11] Sommese, R., Jonker, M., van Rijswijk-Deij, R., Dainotti, A., Claffy, K., and Sperotto, A., "The Forgotten Side of DNS: Orphan and Abandoned Records," *Proceedings of the 2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, online, September 2020, pp. 538-543, doi: 10.1109/EuroSPW51379.2020.00079.

[12] Akiwate, G., Jonker, M., Sommese, R., Foster, I., Voelker, G.M., Savage, E., and Claffy K., "Unresolved Issues: Prevalence, Persistence, and Perils of Lame Delegations," *Proceedings of the ACM Internet Measurement Conference 2020*, online, October 27-29, 2020, pp. 281-294. https://doi.org/10.1145/3419394.3423623

[13] Sommese, R., Bertholdo, L., Akiwate, G., Jonker, M., van Rijswijk-Deij, R., Dainotti, A., Claffy, K.C. and Sperotto, A., "Manycast2: Using anycast to measure anycast," *Proceedings of the ACM Internet Measurement Conference2020*, online, October 27-29, 2020, pp. 456-463.

[14] University of Twente, "Anycast Census dataset", URL: *https://github.com/ut-dacs/Anycast-Census* Accessed 15 February 2022.

[15] Sommese, R., Akiwate, G., Jonker, M., Moura, G. C., Davids, M., van Rijswijk-Deij, R., Sperotto, A. (2021, September). "Characterization of Anycast Adoption in the DNS Authoritative Infrastructure," *Proceedings of the Network Traffic Measurement and Analysis Conference (TMA'21),* online, September 14-15, 2021.

[16] University of Twente and CAIDA, UCSD, "MADDVIRP Architecture Prototype," URL: *https://maddvipr.org/deadline/documents/maddvipr.pdf* October 31, 2021, Accessed 15 February 2022.

[17] Galvin, J., "Afilias to Protect TLDs Against Potential "Orphan Glue" Exploits", Circleid, August 11, 2020,
URL: *https://circleid.com/posts/20200811-afilias-to-protect-tlds-against-potential-orphan-glue-exploits* Accessed 15 February 2022.

[18] Sommese, R., "MAnycast^2: Using anycast to measure anycast," *Apnic blog,* Dec 15, 2020, URL: https://blog.apnic.net/2020/12/15/manycast2-using-anycast-to-measure-anycast/ Accessed 15 February 2022.

APPENDIX A Publications and Presentations

Publications:

Sommese, R., Moura, G., Jonker, M., Rijswijk-Deij, R. V., Dainotti, A., Claffy, K. C., & Sperotto, A (2020) "When Parents and Children Disagree: Diving into DNS Delegation Inconsistency," In Sperotto A., Dainotti A., and Stiller B., "Passive and Active Measurement," *Proceedings of Passive and Active Measurement Conference 2020,* online, March 2020, *Lecture Notes in Computer Science*, vol 12048, Springer, Cham. https://doi.org/10.1007/978-3-030-44081-7_11

Sommese, R., Jonker, M., van Rijswijk-Deij, R., Dainotti, A., Claffy, K., and Sperotto, A., "The Forgotten Side of DNS: Orphan and Abandoned Records," *Proceedings of the 2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, online, September 2020, pp. 538-543, doi: 10.1109/EuroSPW51379.2020.00079.

Akiwate, G., Jonker, M., Sommese, R., Foster, I., Voelker, G.M., Savage, E., and Claffy K., "Unresolved Issues: Prevalence, Persistence, and Perils of Lame Delegations," *Proceedings of the ACM Internet Measurement Conference 2020*, online, October 27-29, 2020, pp. 281-294. https://doi.org/10.1145/3419394.3423623

Sommese, R., Bertholdo, L., Akiwate, G., Jonker, M., van Rijswijk-Deij, R., Dainotti, A., Claffy, K.C. and Sperotto, A., "Manycast2: Using anycast to measure anycast," *Proceedings of the ACM Internet Measurement Conference2020*, online, October 27-29, 2020, pp. 456-463.

Sommese, R., Akiwate, G., Jonker, M., Moura, G. C., Davids, M., van Rijswijk-Deij, R., Sperotto, A. (2021, September). "Characterization of Anycast Adoption in the DNS Authoritative Infrastructure," *Proceedings of the Network Traffic Measurement and Analysis Conference (TMA'21),* online, September 14-15, 2021.

Presentations

Akiwate, G., "Unresolved Issues: Prevalence, Persistence, and Perils of Lame Delegations", *Workshop on Active Internet Measurements: Knowledge of Internet Structure: Measurement, Epistemology, and Technology (AIMS-KISMET)*, La Jolla, CA, US, February 26-28, 2020, URL: https://www.caida.org/catalog/media/2020_unresolved_issues_imc/unresolved_issues_imc.pdf Accessed 15 February 2022.

Sommese, R., "When Parents and Children Disagree: Diving into DNS Delegation Inconsistency," *Passive and Active Measurement Conference 2020*, online, March 2020, URL: https://www.caida.org/catalog/media/2020_when_parents_children_disagree_pam/when_parents_children_disagree_pam.pdf Accessed 15 February 2022

Sommese, R., "When Parents and Children Disagree: Diving into DNS Delegation Inconsistency (RACI)," RIPE80, online, 12-14 May 2020, URL: https://ripe80.ripe.net/archives/video/324/

Sommese, R., "Manycast2: Using anycast to measure anycast," *ACM Internet Measurement Conference2020*, online, October 27-29, 2020, URL: https://www.caida.org/catalog/media/2020_manycast2_imc/manycast2_imc.pdf

Sommese, R., "Characterization of Anycast Adoption in the DNS Authoritative Infrastructure," *Network Traffic Measurement and Analysis Conference (TMA'21),* online, September 14-15, 2021, URL: https://indico.dns-oarc.net/event/40/contributions/870/attachments/851/1545/OARC-Anycast.pdf

Sommese, R., "Characterization of Anycast Adoption in the DNS Authoritative Infrastructure," *OARC 36,* online, November 29, 2021, URL: https://indico.dns-oarc.net/event/40/contributions/870/attachments/851/1545/OARC-Anycast.pdf

Sommese, R., "The MADDVIPR Project One Year Later," *NWO-DHS PI Meeting,* online, October 24, 2019, https://maddvipr.org/DHS.pdf

Sommese, R., "Creating a "long-term memory" for the global DNS," *KISMET 2020*, La Jola, February 26, 2020, https://maddvipr.org/Kismet.pdf

Sommese, R., "The Forgotten Side of DNS: Orphan and Abandoned Records," *WTMC 2020*, online, September 7, 2020, https://indico.dns-oarc.net/event/34/contributions/794/attachments/762/1292/OARC33.pdf

Sommese, R., "The Forgotten Side of DNS: Orphan and Abandoned Records," *OARC33,* online, September 28, 2020, https://indico.dns-oarc.net/event/34/contributions/794/attachments/762/1292/OARC33.pdf Accessed 15 February 2022.

Jonker, M., Sommese, R., "DNSAttackStream: Impact of DDoS attacks against DNS Infrastructure," *Technical Workshop,* July 14, 2021, https://maddvipr.org/DNSAttackStream-DUST.pdf

Sommese, R., Jonker, M., "DNS Transparency Logs," *DINR 2021,* online, November 17, 2021, *https://maddvipr.org/DTL.pdf*

Sommese, R., "Characterization of Anycast Adoption in the DNS Authoritative Infrastructure," November 29, 2021, *Technical Workshop, https://indico.dns-oarc.net/event/40/contributions/870/attachments/851/1545/OARC-Anycast.pdf* Accessed 15 February 2022.

LIST OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS

DNS Domain Name System
DNS-OARC DNS Operations Analysis and Research Center
DNSSEC Domain Name System Security Extensions
DDoS Distributed Denial of Service
DoS Denial of Service
DZDB DNS Zone Database
gTLD generic Top-Level Domains
IETF Internet Engineering Task Force
IP Internet Protocol
NS Name Server
RIPE-NCC Rseaux IP Europens Network Coordination Centre
RSDoS Reflected Spoofed Denial of Service
RRSET Set of all Resource Records
SLD Second Level Domain
SPoF Single Point of Failure
TLD Top-Level Domains
TTL Time to Live
UDP User Datagram Protocol (UDP)