

Unveiling the Weak Links: Exploring DNS Infrastructure Vulnerabilities and Fortifying Defenses

Yevheniya Nosyk, Olivier Hureau, Simon Fernandez, Andrzej Duda, and Maciej Korczyński
Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG
Grenoble, France
first.last@univ-grenoble-alpes.fr

Abstract—In the past decades, DNS has gradually risen into one of the most important systems on the Internet. Malicious actors have long misused it in reflection and amplification DDoS attacks, but given its criticality, DNS quickly became an attractive attack target itself. There appeared a number of activities that make use of domain names and the DNS protocol to perform illegal actions, collectively referred to as DNS abuse. In this paper, we measure the landscape of DNS infrastructure vulnerabilities across millions of recursive resolvers and authoritative nameservers. We enumerate domain names deploying cache poisoning protection (DNSSEC), email authentication (SPF/DMARC), and resolvers accepting DNS requests from arbitrary clients. We show that DNS infrastructure is not sufficiently protected against cybersecurity threats and propose a set of recommendations to mitigate the existing problems. Conducted in the frame of a European Commission project, our findings will be considered for inclusion in the upcoming European Union legislation on cybersecurity.

1. Introduction

The Domain Name System (DNS) is one of the most important building blocks of the modern Internet. It is highly distributed over a number of authoritative nameservers (that store all the data about individual domain names) and recursive resolvers (that traverse the DNS hierarchy to find responses to user requests). Originally created to provide the mapping between human-readable domain names (e.g., `example.com`) and IP addresses (e.g., `192.0.2.1`), it now comprises many more roles such as email authentication [13], [41], [45], SSH key verification [26], or certification authority authorization [27].

Given its pervasive nature, it comes as no surprise that DNS is an attractive target for malicious actors. It has long been misused in Distributed-Denial-of-Service (DDoS) attacks as an efficient reflector and amplifier [40], [55]. Yet, the DNS itself can also fall victim to DNS abuse, which is referred to as any activity that makes use of domain names or the DNS protocol to carry out harmful or illegal activities. Different types of abuse put the burden of remediation actions onto different actors, such as domain resellers, registrars, registries, hosting, or Internet Service Providers (ISPs), among others.

Fighting DNS abuse requires tremendous effort from all the aforementioned entities. Therefore, the focus has recently shifted to preventing abuse rather than coping with its consequences. In September 2022, ICANN

launched KINDNS [33], [69], an initiative to promote the adoption of best current practices for DNS operators running authoritative nameservers (whether critical or not) and recursive resolvers (whether public or closed). In particular, the participants are required to provide DNSSEC validation to their clients and cryptographically sign their domain names. Furthermore, the updated European Network and Information Security (NIS2) Directive [70] considers DNS service providers as those belonging to highly critical sectors. Consequently, they are required to adopt a set of cybersecurity risk-management measures such as “security in network and information systems acquisition, development and maintenance, including vulnerability handling and disclosure.”

In 2020, The European Commission adopted The Cybersecurity Strategy in the Digital Decade [18] that defines its roadmap to secure the Internet: i) hardening the security of connected services, ii) effective response to cyberattacks, and iii) cooperation with partners around the world. As part of its ongoing efforts to improve cyber resilience in the European Union, we were contracted by the European Commission to perform a study on the domain name system abuse [12]. In this paper, we present three measurements covering DNS infrastructure abuse. We study the deployment of server-side DNSSEC, email authentication mechanisms (SPF/DMARC), as well as the landscape of open DNS resolvers. We show that currently deployed mechanisms do not adequately protect domain owners, resolver operators, and end users from cybersecurity threats and propose a set of recommendations to secure the core DNS infrastructure. The proposed recommendations will be considered for inclusion in the upcoming EU legislation on cybersecurity.

The rest of this paper is organized as follows. Section 2 provides the deployment state of DNSSEC. Section 3 enumerates open DNS resolvers and discusses the potential threats. Section 4 analyzes two email authentication mechanisms. Section 5 discusses ethical considerations of this study while Section 6 presents related work. We conclude the paper in Section 7.

2. DNSSEC Deployment

2.1. Motivation

The Domain Name System (DNS) was originally designed over 30 years ago. As security was not the primary concern at the time, the early DNS standard was found

TABLE 1. TOP 20 TLDs WITH THE HIGHEST NUMBER OF SECOND-LEVEL DOMAINS IN OUR INPUT LIST.

Rank	TLD	Count	Type	Rank	TLD	Count	Type
1.	com	145,475,053	gTLD	11.	tk	2,298,943	ccTLD
2.	net	12,213,558	gTLD	12.	ga	2,249,643	ccTLD
3.	de	9,601,890	ccTLD	13.	fr	2,098,489	ccTLD
4.	org	9,540,343	gTLD	14.	cn	1,949,840	ccTLD
5.	uk	4,263,606	ccTLD	15.	it	1,758,075	ccTLD
6.	info	3,492,481	gTLD	16.	ml	1,657,468	ccTLD
7.	ru	3,473,332	ccTLD	17.	eu	1,559,517	ccTLD
8.	nl	2,741,787	ccTLD	18.	au	1,557,872	ccTLD
9.	xyz	2,516,448	gTLD	19.	cf	1,487,356	ccTLD
10.	br	2,309,677	ccTLD	20.	online	1,443,770	gTLD

vulnerable to many classes of attacks [5]. One of them is *cache poisoning*: when a malicious actor sends a forged reply to a recursive resolver before the genuine reply from an authoritative resolver arrives, it stays in the recursive resolver cache. Such a specifically crafted packet can redirect genuine clients to bogus websites, mail or name servers.

DNS security extensions (DNSSEC) solve the problem by introducing origin authentication and data integrity [3], [64], [65] using the public key infrastructure. However, DNSSEC is only effective when deployed universally. We analyzed 251 million domain names and found that a small fraction of them attempted to deploy DNSSEC. Even fewer were correctly signed. We further show that while DNSSEC helps secure certain aspects of DNS, it is also prone to new types of attacks and should be implemented with great caution.

2.2. Background

DNSSEC modifies the normal DNS operation by introducing two new concepts: zone signing and response validation. Zone owners generate public/private key pairs. Private keys are used to sign resource record sets (RRsets) and produce RRSIG signatures. The corresponding DNSKEY public keys verify the signatures. Although not required by the DNSSEC standard, there are usually two key pairs - the Key Signing Key pair (KSK) and the Zone Signing Key pair (ZSK). KSK only signs DNSKEY RRset and its digest is published in the parent zone as the DS resource record. ZSK signs the remaining RRsets.

Zone signing does not protect from manipulation if the keys and signatures are not cryptographically verified. DNSSEC-validating recursive resolvers are pre-configured with one trust anchor, usually the root zone public key (or its digest). The validator follows the chain of DS-DNSKEY resource records from the root zone down the domain name tree to the requested domain name. It ensures that the digests correspond to the public keys and that the public keys verify the signatures. If all the checks are successful, it returns the response with NOERROR status code and SERVFAIL otherwise.

2.3. Measurements

We analyze DNSSEC deployment at two different levels. We first show that the majority of TLDs are signed and can be used to publish DS records of their children. We then switch to second-level domain names and observe that DNSSEC suffers low deployment rates.

2.3.1. Top-Level Domains. The operators of DNSSEC-signed zones assume that validating recursive resolvers will be able to establish a chain of trust from the trust anchor down to the zone. Since 2010, such a universally accepted trust anchor is the root zone KSK [30]. Once the root zone was signed, TLD operators had an opportunity to sign their zones and upload DS records to the root zone. ICANN publishes a daily report on the DNSSEC adoption at the TLD level. As of July 2021, 1,372 out of 1,498 TLDs are signed and publish a key hash at the root [35]. The last generic top-level domain was signed in December 2020 [32] and all the 126 unsigned TLDs are country-codes. Note that to implement DNSSEC, the TLD operator must sign the TLD zone. It is the first and most critical step in implementing DNSSEC. As one of the safeguards proposed by ICANN, all operators of new gTLDs are required to sign the TLD zone [31].

Recommendation: *Similarly to gTLD registries, the registry operators of ccTLDs should be required to sign TLD zone files with DNS security extensions (DNSSEC) and facilitate its deployment according to good practices.*

2.3.2. Second-Level Domains. DNSSEC-signed zones are different from the unsigned ones as they publish additional resource records: DS, DNSKEY, RRSIG, and NSEC(3) that can be queried by recursive resolvers as any other regular resource records such as A, NS, etc. We rely on this fact to enumerate second-level domains that attempted to deploy DNSSEC. We use zdns [39] scanner to send DS and DNSKEY requests efficiently at scale. We operate it in the nameserver mode so that it forwards all the requests to the recursive resolver of our choice. We then set up a resolver using BIND9 [38]. By default, it performs validation of all the received responses. However, we disable this function so that we receive the responses even if they are bogus. While scanning for DNSKEYs, we capture all the incoming traffic and extract RRSIG signatures returned along with DNSKEYs. At this stage, we only check for the presence of resource records and not their validity.

We analyzed the DNSSEC deployment of more than 251 million second-level domain names, representing 1,376 TLDs (Table 1 shows the top 20 TLDs by the number of domains). We collected .com, .net, .org, .biz, .tel, .info legacy gTLDs, and new gTLDs made available to us by the ICANN Centralized Zone Data Service (CZDS) [34], as well as .se and .nu ccTLD zone files [36]. We also developed the scanning platform to crawl all the websites of known domains to retrieve newly observed domain names. Note that for some TLDs for which we have access to their zone files, we evaluate the DNSSEC deployment for all domain names. However, for most ccTLDs, we assess the deployment based on all enumerated domains rather than all registered domain names (e.g., 9.6 million .de domain names, 3.5 million .ru domain names, or 2.7 .nl domain names). Therefore, the results represent the approximate rates of DNSSEC deployment per TLD.

Overall, 227 million domain names returned NOERROR responses to our scanner for both (DS and DNSKEY) queries. We refer to them as *responsive* domains. We exclude the remaining 24 million domains from the further analysis, as we cannot determine whether

TABLE 2. TOP 20 TLDs WITH THE HIGHEST NUMBER OF SECOND-LEVEL DOMAINS FALLING INTO EACH CATEGORY. THE RATIO IS COMPUTED FOR ALL THE RESPONSIVE DOMAINS.

Rank	Unsigned			Incorrectly Signed			Correctly Signed		
	TLD	Count	Ratio (%)	TLD	Count	Ratio (%)	TLD	Count	Ratio (%)
1.	com	122,236,139	93.85	com	4,905,793	3.77	com	3,105,826	2.38
2.	net	10,403,214	96.30	ru	203,715	6.43	nl	1,367,067	51.43
3.	de	9,230,789	97.84	nl	162,992	6.13	se	676,318	54.84
4.	org	8,309,362	96.33	net	95,869	0.89	cz	418,299	59.23
5.	uk	3,939,182	96.88	org	52,810	0.61	net	303,482	2.81
6.	info	2,970,680	97.20	se	43,019	3.49	fr	292,072	14.45
7.	ru	2,959,669	93.44	eu	41,999	2.81	pl	279,901	21.31
8.	ga	2,213,920	99.95	fr	35,661	1.76	br	265,991	12.42
9.	tk	2,196,953	99.88	de	31,482	0.33	org	263,955	3.06
10.	br	1,866,791	87.20	cz	30,120	4.26	eu	237,625	15.87
11.	xyz	1,844,580	97.92	be	28,377	3.47	be	208,268	25.47
12.	fr	1,693,420	83.78	pl	24,541	1.87	dk	200,016	29.70
13.	it	1,685,931	99.10	uk	21,705	0.53	de	172,621	1.83
14.	ml	1,644,070	99.92	xn-plai	21,630	6.94	no	151,435	48.82
15.	cn	1,636,199	99.91	co	17,798	1.60	sk	105,044	47.74
16.	au	1,499,598	99.56	info	15,855	0.52	uk	104,962	2.58
17.	cf	1,480,148	99.95	nu	12,690	7.94	ch	93,150	7.42
18.	gq	1,272,683	99.98	no	12,614	4.07	nu	81,041	50.68
19.	ca	1,269,112	98.33	hu	11,903	3.00	hu	71,959	18.13
20.	eu	1,217,662	81.32	it	11,828	0.70	info	69,699	2.28

TABLE 3. TOP 20 GENERIC TLDs WITH THE HIGHEST NUMBER OF SECOND-LEVEL DOMAINS FALLING INTO EACH CATEGORY. THE RATIO IS COMPUTED FOR ALL THE RESPONSIVE DOMAINS.

Rank	Unsigned			Incorrectly Signed			Correctly Signed		
	TLD	Count	Ratio (%)	TLD	Count	Ratio (%)	TLD	Count	Ratio (%)
1.	com	122,236,139	93.85	com	4,905,793	3.77	com	3,105,826	2.38
2.	net	10,403,214	96.30	net	95,869	0.89	net	303,482	2.81
3.	org	8,309,362	96.33	org	52,810	0.61	org	263,955	3.06
4.	info	2,970,680	97.20	info	15,855	0.52	info	69,699	2.28
5.	xyz	1,844,580	97.92	online	10,643	0.94	app	58,808	10.45
6.	online	1,101,814	97.30	xyz	8,840	0.47	page	52,368	67.34
7.	club	748,133	98.56	shop	5,274	1.02	dev	48,187	24.27
8.	vip	510,011	99.84	site	5,131	1.06	xyz	30,391	1.61
9.	shop	501,711	97.50	dev	3,719	1.87	online	19,985	1.76
10.	app	500,579	88.94	app	3,434	0.61	ovh	15,413	37.29
11.	site	474,974	98.09	store	2,860	0.91	one	12,742	25.57
12.	top	411,033	99.45	tech	2,568	1.28	realty	9,272	79.63
13.	icu	398,970	99.80	club	2,480	0.33	club	8,453	1.11
14.	store	306,033	97.45	cloud	1,508	1.10	tech	8,048	4.00
15.	live	294,922	98.35	mobi	1,449	0.65	shop	7,611	1.48
16.	work	266,555	99.36	space	1,368	0.82	cloud	6,461	4.70
17.	mobi	220,399	98.79	top	1,040	0.25	store	5,139	1.64
18.	tech	190,506	94.72	website	1,038	0.79	studio	4,460	8.85
19.	space	161,307	97.21	xn-placf	941	6.81	live	4,203	1.40
20.	dev	146,628	73.86	fun	886	0.87	site	4,112	0.85

they do not publish some of the resource records or we could not retrieve them for other reasons (temporary network failures, etc.).

We first check how many responsive domains contain one or more DNSSEC resource records: DNSKEY, RRSIG, and/or DS. The presence of such records does not necessarily mean that domains are *correctly* signed, but rather signifies that domain owners attempted to do so. Only 6.7% (15.2 million) of responsive domains publish at least one DNSSEC resource record. Half of them fail to provide all three RRs. Such misconfigurations have different consequences:

- DNSKEY-RRSIG, DNSKEY, RRSIG: the lack of DS is a common misconfiguration, as this record needs to be manually added to the parent zone (through the registrar control panel). It was previously shown that around 30% of domains that publish DNSKEY do not have an associated DS [8]. The responses from these domains are considered *insecure* by the DNSSEC standard [3]. They will

not fail the validation check by recursive resolvers, but without a complete chain of trust, we cannot conclude whether the domain is correctly signed. Such DNS zones are referred to as islands of security and can only be used to validate their child zones (if recursive resolvers trust their keys). There are 5.7 million second-level domains from 748 TLDs that fail to provide the DS record while providing the two others (DNSKEY and/or RRSIG).

- DNSKEY-DS, RRSIG-DS, DS: the domains with the DS records at the delegation point have the complete chain of trust and will be verified by validating recursive resolvers. Because of the missing signatures (RRSIG) and/or public keys (DNSKEY), the validation will fail (the responses from such domains are called *bogus*), and the end clients will receive SERVFAIL in response to their requests. Such misconfigurations, combined with using validating resolvers, effectively make these domains

TABLE 4. TOP 20 COUNTRY-CODE TLDs WITH THE HIGHEST NUMBER OF SECOND-LEVEL DOMAINS FALLING INTO EACH CATEGORY. THE RATIO IS COMPUTED FOR ALL THE RESPONSIVE DOMAINS.

Unsigned			Incorrectly Signed			Correctly Signed		
TLD	Count	Ratio (%)	TLD	Count	Ratio (%)	TLD	Count	Ratio (%)
de	9,230,789	97.84	ru	203,715	6.43	nl	1,367,067	51.43
uk	3,939,182	96.88	nl	162,992	6.13	se	676,318	54.84
ru	2,959,669	93.44	se	43,019	3.49	cz	418,299	59.23
ga	2,213,920	99.95	eu	41,999	2.81	fr	292,072	14.45
tk	2,196,953	99.88	fr	35,661	1.76	pl	279,901	21.31
br	1,866,791	87.20	de	31,482	0.33	br	265,991	12.42
fr	1,693,420	83.78	cz	30,120	4.26	eu	237,625	15.87
it	1,685,931	99.10	be	28,377	3.47	be	208,268	25.47
ml	1,644,070	99.92	pl	24,541	1.87	dk	200,016	29.70
cn	1,636,199	99.91	uk	21,705	0.53	de	172,621	1.83
au	1,499,598	99.56	xn-plai	21,630	6.94	no	151,435	48.82
cf	1,480,148	99.95	co	17,798	1.60	sk	105,044	47.74
gq	1,272,683	99.98	nu	12,690	7.94	uk	104,962	2.58
ca	1,269,112	98.33	no	12,614	4.07	ch	93,150	7.42
eu	1,217,662	81.32	hu	11,903	3.00	nu	81,041	50.68
ch	1,158,848	92.27	it	11,828	0.70	hu	71,959	18.13
us	1,142,047	97.32	dk	11,368	1.69	co	46,556	4.20
nl	1,127,965	42.44	us	10,369	0.88	us	21,048	1.79
co	1,044,857	94.20	sk	8,324	3.78	ca	16,747	1.30
pl	1,009,203	76.82	br	8,064	0.38	io	15,646	6.10

unreachable. There are 112,648 second-level domains from 422 TLDs that fail to provide DNSKEY and/or RRSIG while providing DS record.

These preliminary findings are alarming. The great majority of tested domains do not contain any resource records that would signal the willingness of domain owners to deploy DNSSEC. Only 15.2 million domains contain one or more DNSSEC-related resource records (DNSKEY, DS, RRSIG). However, we see straight away that 37.6% of them are, in the best case, islands of security (because of missing DS), and 0.7% of them will fail the validation (because of missing public keys and/or signatures).

Note that in addition to TLD registries, registrars also play a key role in the implementation of DNSSEC, as they must add the DS record to the parent zone maintained by the TLD registry. The lack of support from registrars means that all domain names managed by these registrars cannot be signed. The Danish Ministry of Business has implemented a law requiring the .dk registry to ensure that all registrars that offer domain names in the .dk domain support DNSSEC no later than January 1, 2021, and offer DNSSEC signing to registrants [53]. Some registrars not only facilitate the addition of a DS record to a master zone but provide the “one-click” DNSSEC deployment as a paid option (e.g., GoDaddy) or even at no cost (e.g., OVH SAS). The second option is one of the best ways to increase the DNSSEC deployment on a massive scale.

Recommendation: *To facilitate the deployment of DNSSEC, domain administrators (registrants) should have easy access to DNSSEC signing of domain names within the TLD. TLD registries should require all registrars that offer domain names in the TLD to support DNSSEC signing for registrants.*

The domains that do provide all three resource records (9.4 million) are likely to be correctly signed but need further validation. We switch our BIND9 recursive resolver into validating mode and query these domains for the SOA and DNSKEY records. The validating recursive resolver retrieves the requested resource records, performs additional queries to establish the chain of trust, and validates

the signatures. The results are reassuring: 98.1% of domains publishing all three resource records correctly sign both DNSKEY and SOA resource records. Thus, we can conclude that the presence of all the necessary DNSSEC resource records results in a high chance that the zone is *correctly* signed.

Based on our measurements, we categorize all the responsive domains (227 million) into three groups:

- *Unsigned* (212 million): the domains that do not publish any DNSSEC resource records (DNSKEY, DS and RRSIG).
- *Incorrectly signed* (6 million): the domains that either publish some of DNSSEC resource records or all of them, but fail to validate.
- *Correctly signed* (9 million): the domains that publish all the DNSSEC records and, when queried by a validating resolver, provide correctly signed responses.

Tables 2, 3, 4, and 5 provide the information on what TLDs have the highest numbers of second-level domains falling into each category. Table 2 displays top 20 TLDs of unsigned, incorrectly signed, and correctly signed domains. Tables 3 and 4 show similar ranking among generic and country-code TLDs. Table 5 shows the number and the DNSSEC-deployment rate of European Union TLDs in each category.

As mentioned earlier, we computed the rates for most ccTLDs based on a large sample of identified domain names because we do not have access to the zone files and the complete list of domain names. Therefore, the presented rates provide an approximation of the actual adoption. The DNSSEC adoption rates are not different from the general population and are rather modest—21 out of 34 TLDs consist of more than 90% of unsigned domains. On the contrary, the .cz TLD exhibits the highest proportion of correctly signed second-level domains. The cz.nic domain registry achieved it thanks to incentivizing registrars and ISPs economically and supporting them technically [21]. Moreover, DNSSEC is a part of the governmental initiative called “Digital Czech Republic

v. 2.0” [52]. Swedish country-code TLD comes second with the majority (54.84%) of correctly signed domains. The `.se` registry provides guidance on the DNSSEC deployment [37] and price incentives. The Dutch TLD `.nl` has high DNSSEC adoption rates (51.43%) thanks to the support from both the government and SIDN, the registry of `.nl` domains [67]. Registrars are charged lower fees for DNSSEC-signed domain names than for unsigned domain names. Finally, OVH (registrar for multiple TLDs and registry for `.ovh` domain names) proposes free DNSSEC signing to all its customers [59] with “one-click” regardless of the domain TLD, which resulted in a high DNSSEC adoption rate of 37% for the `.ovh` domain names as shown in Table 2.

The examples of the `.cz`, `.se`, `.no`, or `.nl` TLDs show that price incentives are the main driving factor behind the deployment of DNSSEC. All these registry operators are among those that have used such schemes.

Recommendation: *As an incentive to the deployment of DNSSEC, TLD registries might offer discounts for DNSSEC-signed domain names.*

2.4. Challenges

DNSSEC has technically solved the problem of forged DNS replies. However, the administrators of signed zones face additional maintenance issues such as key management and signature expiration. We discuss DNSSEC challenges in the remainder of this section.

2.4.1. Amplification of DDoS Attacks. DNS has long been known as one of the most used protocols to launch reflection and, especially, amplification DDoS attacks [2], [47], [66]. DNSSEC introduced a non-negligible overhead to the normal DNS operation because signed responses are larger in size. Van Rijswijk-Deij et al. [74] analyzed 2.5 million signed domains and a sample of unsigned domains across 6 TLDs and their amplification factors. While regular queries (A, AAAA, DNSKEY, NSEC3, MX, NS, TXT) do increase the amplification factor compared to normal DNS, it mostly does not exceed the theoretical upper bound. A more serious amplifier is ANY type query, which results in the amplification factor of 47.2 for signed domains versus 5.9 for unsigned. Zone administrators cannot prevent attackers from querying their nameservers. Yet, they can block or provide minimal responses to ANY queries [1] and configure the nameservers with response rate/size limiting.

2.4.2. Signature Validity. RRSIG signatures introduce the notion of absolute time in DNS. The two fields (Signature Inception and Signature Expiration) are timestamps that specify the time period during which the signature can be used for validation. Validating recursive resolvers use “their own notion of current time” [65] to check that the signature expiration field is greater than or equal to it. We examined 12.8 million signatures across 10.6 million second-level domains and found that 17,376 of them are expired. Responses with such signatures are *bogus*. Zone administrators should make sure that their signatures are always valid. RFC 6781 lists more time-related considerations in DNSSEC [42]. For example, signed zones are advised to have TTL values

smaller than the signature validity period, which will avoid data being flushed from recursive resolvers caches once signature expiration time is reached.

2.4.3. Key Management. For DNSSEC to be cryptographically secure, zone administrators should only sign their zones with recommended algorithms defined in RFC 8624 [77]. We checked whether the domains in our dataset publishing DNSKEY records (15.1 million) adhere to this standard. We found that 25.9% of all the DNSKEY (25.4 million) implement not recommended algorithms. Only few domains (507) implement algorithms that *must not* be used.

Chung et al. [8] closely examined some of the common issues when it comes to key management in DNSSEC. Key reuse occurs when one private key is used to sign multiple domains. Although it was found that only 0.5% of examined keys are shared, one KSK and ZSK were shared among 130,000 domains. If a private key gets compromised, these domains will be affected at once. Another concern is the key size. The DNSSEC standard does not dictate the key size requirements but the authors refer to NIST recommendations [6]. They found that 91.7% of examined ZSKs were not meeting the minimal key size requirements.

2.4.4. DNSSEC Validation. To protect end-users from cache poisoning attacks, local resolvers must verify the chain of trust to ensure the integrity and authenticity of domain name resolutions. Even the complete deployment of DNSSEC by TLD registries, registrars, and registrants will not protect end users if DNS resolvers do not perform validation. One of the challenges is to measure whether ISPs perform validation, as it requires performing DNS queries from within the tested networks. In addition, it is challenging to measure the impact of the DNSSEC deployment on global security because the detection of cache poisoning attacks can generally be done at the ISP level or using passive DNS data.

Recommendation: *Internet Service Providers that operate DNS resolvers should configure DNSSEC validation to protect end users from cache poisoning attacks and ensure the integrity and authenticity of domain name resolutions.*

2.5. Discussion

DNSSEC remains the most effective way to fight DNS cache poisoning but only when universally deployed. Surprisingly, 126 TLDs are still not signed. Consequently, their child zones cannot fully deploy DNSSEC because they will not have the complete chain of trust. Out of analyzed 227 million active second-level domain names, a small fraction (9.2 million) are correctly signed.

The DNSSEC operation is complex and involves multiple parties: registrants, zone administrators (if different from registrants), registrars, TLDs, and operators of recursive DNS resolvers. To increase the adoption (and validation) of DNSSEC, everyone needs to participate. The remaining unsigned country-code TLDs should adopt DNSSEC to improve their reputation and enable their customers to sign their domains. They should also incentivize registrars to deploy it. Registrars, on their side, can

TABLE 5. COUNTRY-CODE TLDs OF EUROPEAN UNION MEMBERS WITH THE HIGHEST NUMBER OF SECOND-LEVEL DOMAINS FALLING INTO EACH CATEGORY. THE RATIO IS COMPUTED FOR ALL THE RESPONSIVE DOMAINS.

Rank	Unsigned			Incorrectly Signed			Correctly Signed		
	TLD	Count	Ratio (%)	TLD	Count	Ratio (%)	TLD	Count	Ratio (%)
1.	xn-qxa6a	8	100.00	nl	162,992	6.13	cz	418,299	59.23
2.	xn-qxam	321	99.69	cz	30,120	4.26	se	676,318	54.84
3.	ie	115,817	99.65	xn-e1a4c	4	3.96	nl	1,367,067	51.43
4.	xn-90ae	475	99.58	sk	8,324	3.78	sk	105,044	47.74
5.	hr	46,763	99.54	se	43,019	3.49	dk	200,016	29.70
6.	lt	94,003	99.42	be	28,377	3.47	be	208,268	25.47
7.	mt	4,116	99.18	hu	11,903	3.00	pl	279,901	21.31
8.	si	66,691	99.15	eu	41,999	2.81	hu	71,959	18.13
9.	it	1,685,931	99.10	lv	1,392	2.67	eu	237,625	15.87
10.	gr	211,576	99.10	gl	44	1.90	fr	292,072	14.45
11.	cy	5,508	98.80	pl	24,541	1.87	ee	8,599	13.09
12.	bg	33,108	98.75	fr	35,661	1.76	xn-e1a4c	13	12.87
13.	at	663,527	97.95	dk	11,368	1.69	lv	4,425	8.49
14.	ro	301,433	97.94	fo	30	1.59	lu	1,620	4.31
15.	de	9,230,789	97.84	ro	4,753	1.54	pt	4,692	2.91
16.	es	626,866	97.43	pt	2,069	1.28	fi	6,712	2.60
17.	fo	1,831	97.14	cy	65	1.17	gl	49	2.12
18.	fi	248,835	96.47	fi	2,404	0.93	es	13,302	2.07
19.	gl	2,220	95.98	lu	345	0.92	de	172,621	1.83
20.	pt	154,656	95.81	mt	34	0.82	at	11,218	1.66
21.	lu	35,633	94.77	bg	260	0.78	fo	24	1.27
22.	lv	46,301	88.84	it	11,828	0.70	ro	1,572	0.51
23.	ee	56,769	86.42	gr	1,398	0.65	bg	158	0.47
24.	fr	1,693,420	83.78	si	405	0.60	xn-90ae	2	0.42
25.	xn-e1a4c	84	83.17	es	3,202	0.50	lt	242	0.26
26.	eu	1,217,662	81.32	ee	321	0.49	gr	527	0.25
27.	hu	313,065	78.87	at	2,638	0.39	si	169	0.25
28.	pl	1,009,203	76.82	hr	159	0.34	it	3,527	0.21
29.	be	581,024	71.06	de	31,482	0.33	hr	58	0.12
30.	dk	462,170	68.62	lt	305	0.32	ie	128	0.11
31.	sk	106,687	48.48	xn-qxam	1	0.31	cy	2	0.04
31.	nl	1,127,965	42.44	ie	274	0.24	-	-	-
33.	se	513,880	41.67	-	-	-	-	-	-
34.	cz	257,861	36.51	-	-	-	-	-	-

TABLE 6. TOP 20 AUTONOMOUS SYSTEMS BY THE NUMBER OF OPEN RESOLVERS

Rank	ASN	Organization	IPv4 Resolvers	ASN	Organization	IPv6 Resolvers
1.	4134	China Telecom	260,649	6939	Hurricane Electric	4,458
2.	4837	China Unicom	189,714	63949	Linode	548
3.	45090	Tencent-CN	107,769	3462	HiNet	415
4.	4766	Korea Telecom	67,557	4837	China Unicom	364
5.	47331	TTNET A.S.	58,693	8966	Etisalat-AS	351
6.	5617	Orange Polska	52,568	12322	Free SAS	332
7.	3462	HiNet	36,868	4812	China Telecom	294
8.	4812	China Telecom	33,432	1241	Forthnet	286
9.	9318	SK Broadband	26,903	51167	Contabo	228
10.	4808	China Unicom	26,762	27839	Comteco	184
11.	12389	Rostelecom	24,989	16276	OVH	179
12.	209	Centurylink	24,979	7922	Comcast	163
13.	7713	Telekomunikasi Indonesia	21,475	4134	China Telecom	159
14.	4538	China Education and Research Network Center	18,866	37564	Wirulink Pty Ltd	153
15.	9808	China Mobile	17,838	8100	QuadraNet Enterprises LLC	137
16.	58224	Iran Telecommunication Company	16,036	23910	China Next Generation Internet CERNET2	115
17.	45804	Meghbel Cable & Broadband Services	15,624	3303	Swisscom (Schweiz) AG	110
18.	32708	Root Networks	15,502	3356	Level 3 Parent, LLC	104
19.	3269	Telecom Italia S.p.A.	12,371	14061	DigitalOcean, LLC	102
20.	58659	Quest Consultancy	11,918	8251	FreeTel, s.r.o.	102

encourage domain owners to deploy DNSSEC by offering them discounts and facilitating the signing process [9].

3. Open DNS Resolvers

3.1. Motivation

In addition to proactive and reactive actions taken by TLD registries, registrars, hosting providers, or resellers, DNS resolver operators have also an imperative role in securing the DNS infrastructure. Historically, mainly Internet Service Providers (ISPs) were responsible for

maintaining DNS resolvers that resolve domain names on behalf of end users. However, several companies such as Google [25], Cloudflare [11], Quad9 [63], or OpenDNS [10] have been offering free and public DNS servers as an alternative way to connect to the Internet in recent years. One of the main advantages of using public DNS resolvers is to speed up domain name resolution, thereby improving the quality of experience for end users.

Moreover, regardless of whether the DNS resolver service is operated by local ISPs or public resolver operators, they should apply certain measures to improve the security of end users. Service operators may subscribe to

TABLE 7. TOP 20 IPv4 AUTONOMOUS SYSTEMS BY THE RATIO OF OPEN RESOLVERS.

Rank	ASN	Organization	AS Size	Ratio
1.	269113	Uno Telecom LTDA	1,024	99.5 %
2.	268137	Net Sini Fiber Home Telecomunicação LTDA	1,024	99.5 %
3.	136668	Iana Solutions Digital India	512	99.4 %
4.	263108	Opanet Telecomunicacoes LTDA	2,048	99.3 %
5.	267072	Veloz Net Serviços e Comunicações LTDA	768	99.2 %
6.	267007	Turbo Net Telecom Servicos e Vendas de Equipamento	1,024	99.1 %
7.	134929	Orange City Internet Services	2,048	99.0 %
8.	208070	TILYTEL B., S.L.	1,024	99.0 %
9.	270404	Qualidade Digital Internet e Telecomunicações	1,024	99.0 %
10.	134924	Aph Networks	512	99.0 %
11.	269563	MAX3 TELECOM LTDA	1,024	98.93 %
12.	271003	MARILETE PEREIRA DOS SANTOS	1,024	98.83 %
13.	270657	FNET TELECOM	1,024	98.83 %
14.	34939	NextDNS	768	98.83 %
15.	137045	Athoy Cyber Net	512	98.83 %
16.	47849	Global Communication Net Plc	3,072	98.73 %
17.	269012	Click Net Link Informatica e Telecomunicações LTDA	1,024	98.73 %
18.	265276	SPEED_MAAAX BANDA LARGA LTDA - ME	1,024	98.73 %
19.	271070	Ailson Tavares	1,024	98.63 %
20.	47275	Torjon Wieslaw Radka	1,024	98.63 %

blacklists and should not resolve maliciously registered domain names to their IP addresses. A malicious domain name should resolve with NXDOMAIN indicating that the domain name does not exist or should be resolved to the DNS service provider own blocking site instead of the IP address of the requested malicious domain. The Quad9 [63] system uses threat intelligence from more than a dozen leading cybersecurity companies to provide real-time information about which sites contain malware or other threats. If the system detects that a site a user wants to visit is infected, it automatically blocks the user from accessing it. The public resolver operated by Google does not, in principle, perform any blocking [24]. Instead, malicious URLs (and domain names) are blocked by web browsers (e.g., Chrome, Firefox) using Google Safe Browsing.

The problem is raised by open (misconfigured) DNS servers that facilitate amplification reflection Distributed Denial-of-Service (DRDoS) attacks [47], [57], [66], [73], [75], [78]. Open DNS resolvers accept DNS requests from any end host, which can be misused to either target authoritative nameservers by sending an excessive number of incoming requests or, if combined with IP address spoofing, used to redirect responses to victim end-hosts. Therefore, service providers should significantly reduce the number of misconfigured DNS resolvers to increase the barriers to launching DDoS attacks. In the following sections, we actively scan for open DNS resolvers in IPv4 and IPv6 address spaces and analyze their distribution across organizations and countries.

3.2. Methodology

Scanning for open resolvers requires sending DNS requests to end hosts and inspecting the received responses. The response codes (RCODE) defined in RFC 1035 [54] signal whether the DNS server processes incoming requests. If the query resolution is successful, open resolvers send back the responses to end clients along with the NOERROR status code.

We use three following datasets to scan for open resolvers: IPv4 BGP prefixes [71], IPv6 Hitlist Service [23], and IPv6 addresses learned by traversal from IPv4 to re-

solve IPv6-only domains as described by Nosyk et al. [58]. All three datasets contain globally reachable IP addresses that may be operational recursive resolvers. Each end host from the list receives an A request for the unique domain name under our authority. We developed a software tool that allows us efficiently send DNS packets at a large scale [68].

3.3. Results

3.3.1. Scan Results. We performed IPv4 and IPv6 open resolver scans in March 2021. Having tested more than 2.8 billion routable IPv4 addresses and 3.5 million IPv6 addresses, we discovered 3.4 million IPv4 and 18,843 IPv6 open recursive resolvers. Although the mentioned open resolvers returned the NOERROR responses, they are not necessarily *correctly* operating. We closely inspected the answer section of returned packets and found that 18% IPv4 and 15% IPv6 open resolvers returned empty responses. More importantly, 8.4% and 6.6% of resolvers returned *bogus* replies to our A requests. Previous work has shown that this behavior is likely due to censorship, ad redirection, and other doubtful activities [46]. As the majority of such recursive resolvers return custom responses without contacting authoritative nameservers, their use in DDoS attacks is limited. Thus, we exclude them from further analysis and keep the remaining 2.5 million IPv4 and IPv6 resolvers that can potentially be used as reflectors in DDoS attacks.

3.3.2. Autonomous System Distribution. We map the remaining open resolvers to their Autonomous System numbers (ASN) using `pyasn` [4] and check the PeeringDB [62] and AS Rank [7] for the organization names. Table 6 presents the number of open DNS resolvers per autonomous systems. The top 20 IPv4 organizations are dominated by Asian telecommunication operators, while IPv6 autonomous systems also include transit and hosting providers. In total, open resolvers are present in 24,087 IPv4 and 1,607 IPv6 autonomous systems (34.2% and 7.4% of all those in the BGP routing table as of the beginning of March 2021).

The large absolute number of recursive resolvers may not be surprising if they belong to a large autonomous

TABLE 8. TOP 20 COUNTRIES/TERRITORIES BY THE NUMBER OF OPEN RESOLVERS.

Rank	Country	IPv4 Resolvers	Country	IPv6 Resolvers
1.	China	758,083	USA	2,500
2.	Brazil	323,263	Germany	1,323
3.	USA	180,328	China	1,258
4.	India	117,363	France	880
5.	Republic of Korea	116,749	Republic of Korea	708
6.	Russia	97,287	Taiwan	583
7.	Turkey	78,982	Russia	494
8.	Indonesia	75,157	Czech Republic	409
9.	Poland	73,189	Japan	395
10.	Taiwan	42,577	UK	376
11.	Bangladesh	38,061	Brazil	367
12.	Argentina	34,858	United Arab Emirates	354
13.	France	31,720	Greece	342
14.	Italy	28,916	Thailand	310
15.	Ukraine	27,348	Canada	307
16.	Iran	27,343	Iran	295
17.	Japan	24,808	Vietnam	252
18.	Thailand	22,520	India	244
19.	Hong Kong	20,765	Switzerland	242
20.	Bulgaria	19,992	South Africa	239

system. Thus, we compute a ratio of open resolvers to the size of the address space announced by the IPv4 autonomous systems. Table 7 shows the results: none of the organizations from Table 6 is present in Table 7. These small autonomous systems almost entirely consist of open resolvers. In fact, there are 278 IPv4 ASes for which more than half of the address space is occupied by open resolvers.

3.3.3. Geographic Distribution. We map all the open resolvers to countries using the MaxMind database [51]. Overall, open resolvers are present in 230 countries/territories. Table 8 shows the top twenty countries by the number of open IPv4 and IPv6 resolvers. Eleven countries dominate both the IPv4 and IPv6 ranking. More importantly, the top twenty countries contain the majority of all the open resolvers worldwide: 84.9% in IPv4 and 80.4% in IPv6. Table 9 displays the number of open resolvers in European Union (EU) countries only. The top three countries account for more than 50% IPv4 and 66% IPv6 open resolvers in the EU.

Next, we examine the ratio of open resolvers per region in Table 10. The majority of IPv4 resolvers are located in Asia. IPv6 resolvers are not dominated by a single region, as more than 60% of those are shared between Asia and the European Union. Africa, Oceania, and Europe (outside the European Union) represent the smallest share of open resolvers.

3.4. Discussion

Open resolvers pose an important security threat—they are prone to misuse by attackers and should only be operated when necessary. We discovered over 2.5 million correctly resolving IPv4 and IPv6 open resolvers worldwide. We have shown that they are distributed both in terms of organizations and geographic territories. Nevertheless, most of the open resolvers originate from very few autonomous systems and countries.

Kührer et al. fingerprinted 5.4 million open resolvers and concluded that more than 60% of them were routers,

TABLE 9. DISTRIBUTION OF OPEN RESOLVERS IN EUROPEAN UNION COUNTRIES.

Rank	Country	IPv4 Resolvers	Country	IPv6 Resolvers
1.	Poland	73,189	Germany	1,323
2.	France	31,720	France	880
3.	Italy	28,916	Czech Republic	409
4.	Bulgaria	19,992	Greece	342
5.	Germany	18,352	Netherlands	181
6.	Spain	12,400	Hungary	119
7.	Hungary	10,221	Italy	76
8.	Romania	7,766	Romania	74
9.	Czech Republic	7,508	Austria	72
10.	Netherlands	7,165	Lithuania	64
11.	Sweden	5,945	Sweden	59
12.	Greece	4,962	Poland	53
13.	Austria	3,722	Spain	51
14.	Slovakia	3,663	Bulgaria	48
15.	Portugal	3,646	Slovenia	43
16.	Latvia	3,394	Finland	30
17.	Croatia	2,547	Belgium	30
18.	Denmark	1,877	Denmark	26
19.	Finland	1,738	Portugal	18
20.	Belgium	1,734	Ireland	18
21.	Lithuania	1,178	Croatia	18
22.	Ireland	1,145	Cyprus	17
23.	Cyprus	694	Latvia	13
24.	Slovenia	687	Slovakia	10
25.	Estonia	355	Luxembourg	7
26.	Luxembourg	313	Estonia	6
27.	Malta	250	-	-

TABLE 10. RATIO OF OPEN RESOLVERS PER REGION.

Region	Ratio of IPv4 Resolvers	Ratio of IPv6 Resolvers
Africa	2.4 %	2.1 %
Asia	56.4 %	34.9 %
Europe	6.3 %	8.6 %
European Union	10.1 %	26.9 %
North America	7.8 %	19.0 %
Oceania	0.5 %	1.0 %
South America	16.5%	7.5 %

modems, gateways, and embedded devices [46]. We hypothesize that telecommunication operators do not configure customer equipment correctly. If it is the case, then some national telecommunication operators could eliminate a significant number of open resolvers in their countries (e.g., Orange Polska or Telecom Italia).

Note that this problem has been known for years. In 2013, Jared Mauch presented the Open Resolver Project [50] at the NANOG meeting. He uncovered 34 million DNS servers that responded to UDP/53 probes. Despite different initiatives to mitigate the problem, such as Computer Emergency Response Team (CERT) alerts [14], research indicating the scale of the problem [47], [66], and notifications to operators by ShadowServer, or locally by the national German CERT [20], the issue has still not been resolved.

Recommendation: *National CERT teams should subscribe to data sources that identify open DNS resolvers. National governments and Computer Emergency Response Team (CERT) teams should intensify notification efforts to reduce the number of open DNS resolvers (and other open services), which are among the root causes of distributed reflective denial-of-service (DRDoS) attacks.*

TABLE 11. MOST COMMON SYNTACTICALLY WRONG RULES THAT LEAD TO THE PERMERROR RESULT.

Error type	Example	Correct rule	Count
Too many DNS lookups	-	SPF rule must generate less than 10 DNS query	1,638,092
Two or more SPF records found	-	Must set one SPF record for each domain	691,746
Void lookup limit of 2 exceeded	-	Rules with empty responses must be removed	64,914
More than 10 MX records returned	-	Total number of lookups must be less than 10	27,699
Invalid IP4 address	ip4:xxx.xxx.xxx.xx?all	ip4:xxx.xxx.xxx.xx ?all	16,621

TABLE 12. SCAN RESULTS OF THE SPF RULES.

Rank	Status	Count	Ratio (%)
1.	NOERROR	226,793,645	91.816
2.	SERVFAIL	10,616,307	4.297
3.	REFUSED	5,979,033	2.420
4.	NXDOMAIN	2,475,409	1.002
5.	TIMEOUT	696,076	0.281
6.	AUTHFAIL	275,925	0.111
7.	ERROR	169,679	0.068
8.	TEMPORARY	348	0.0001
SPF record		77,487,889	31.370

TABLE 13. RESULTS OF THE CHECK_HOST FUNCTION EMULATION.

Rank	Status	Count	Ratio (%)
1.	Softfail	35,929,956	46.37
2.	Fail	29,049,907	37.5
3.	Neutral	5,866,297	7.58
4.	Permererror	3,207,817	4.14
5.	None	2,543,870	3.28
6.	Temperror	689,680	0.91
7.	Pass	200,362	0.26

4. SPF and DMARC deployment

4.1. Motivation

Email spoofing is defined as sending emails with a forged sender address in a way that it appears as sent from a legitimate user or on behalf of an organization [48]. Business Email Compromise (BEC) is one of the most financially damaging online crimes [19] and email spoofing is one of the common techniques used in BEC.

The Simple Mail Transfer Protocol (SMTP) does not provide a built-in approach to fight email spoofing. Therefore, the deployment of the email security protocols such as the Sender Policy Framework (SPF) [41], DomainKeys Identified Mail (DKIM) [13], and Domain-based Message Authentication, Reporting, and Conformance (DMARC) [45] is the first line of defense against email spoofing and phishing attacks. In this section, we measure the global adoption of DNS-based email security extensions, namely, SPF and DMARC for 251 million domain names in our database, as described in Section 2.3.2. We do not measure DKIM since it needs access to DKIM subdomains (also known as the `selector` tag). They are not publicly available and can only be retrieved from the header of the received emails.

4.2. Methodology

We use the following approach to measure the deployment of SPF and DMARC. We first scan 251 million domain names for SPF records (found inside DNS `TXT` resource records) using `zdns`. Then, for those domains with SPF records (containing the version string `v=spf1`), we emulate the `check_host()` function as defined by RFC 7208 [41] to evaluate the validity and configurations of the records. The next step is to collect the DMARC rules, published in the `TXT` resource records of the `_dmarc` subdomains of the registered domains (e.g., `_dmarc.example.com`). Finally, we evaluate DMARC rules, i.e., the records starting with `v=DMARC1`, to check their strictness in accepting (delivering to the end users) or rejecting incoming forged emails.

4.3. Results

As shown in Table 12, only 31.37% of the domains have an SPF record, which does not necessarily guarantee protection against email spoofing. Table 13 shows the results of the `check_host()` function emulation for the domains with SPF records. All the domains with the SPF `pass` results are open to email spoofing since they let the sender send emails from any IP address. For other SPF results (e.g., `fail`, `softfail`, `permererror`, etc.), the decision is made by the receiver with the help of DMARC rules specified in the `TXT` resource records of the `_dmarc` subdomain. The SPF `permererror` result means that there is a problem in either parsing or recursive querying SPF rules, which usually happens because of setting a syntactically wrong SPF rule or defining too many DNS lookups (recursions) in the SPF rule [49]. Table 11 shows the most common errors related to the domains with the `permererror` result from the `check_host()` function.

Table 14 shows the scan results for DMARC records. The `NXDOMAIN` status means that there is no DMARC subdomain for the domain name. `NOERROR` indicates that the `_dmarc` subdomain exists but only a small fraction (3.3%) of the domain names contain a `TXT` record with a valid DMARC record. However, having correctly configured DMARC does not necessarily guarantee protection against email spoofing. The final decision about the incoming email delivery is up to the receiver based on the `p` tag (policy) of the DMARC record. It specifies the following actions: i) deliver the message, ii) reject the message, or iii) quarantine the message (flag it as suspicious and, for example, place it into a spam folder). Parsing the DMARC record shows that 49.68% of the domain names with the DMARC record have the `p=none` rule, which means they specified no strict action with respect to incoming messages sent from unauthorized servers. 11.20% of the domains have `p=quarantine` (e.g., labeling the incoming message as spam), and 37.14% have `p=reject`, which means rejecting the incoming message with unauthorized sender based on SPF (and DKIM) rules.

TABLE 14. SCAN RESULTS OF THE DMARC RECORDS.

Rank	Status	Count	Ratio (%)
1.	NXDOMAIN	150,842,488	61.212
2.	NOERROR	78,407,747	31.817
3.	SERVFAIL	10,037,232	4.073
4.	REFUSED	6,019,716	2.442
5.	TIMEOUT	656,653	0.266
6.	AUTHFAIL	284,939	0.115
7.	ERROR	166,427	0.067
8.	TEMPORARY	10,795	0.004
Valid DMARC record		8,129,795	3.299

4.4. Discussion

SPF and DMARC protocols are critical for preventing email spoofing and essential in preventing Business Email Compromise (BEC) fraud, which according to the recent FBI report, caused more than US \$1.8 billion in losses to businesses and individuals in 2020 [19]. Note that securing domain names with SPF and DMARC does not solve the problem of BEC scams, as criminals can register, e.g., misspelled (e.g., using special characters), or internationalized domain names. However, if SPF and DMARC rules are not correctly configured, a cybercriminal can send emails on behalf of target brand domain names, making recipients unable to distinguish legitimate email senders from fraudulent ones. Correctly implemented and strict SPF and DMARC rules can mitigate the problem of domain name spoofing, assuming that recipient mail servers verify and filter emails based on SPF and DMARC rules.

Recommendation: *Security community should intensify efforts to measure the adoption of the SPF and DMARC protocols continuously, especially for high-risk domain names, and raise awareness of the domain spoofing problem among domain owners and email service providers. Correct and strict SPF and DMARC rules can mitigate email spoofing and provide the first line of defense against Business Email Compromise (BEC) scams.*

5. Ethical Considerations

Measurement studies must follow certain precautions so that results are obtained with minimum intervention for tested systems. We use some of the best practices introduced by the measurement community as our guidelines [15], [17], [61]. We set up a simple web page on all the scanner machines with a basic description of our activities and contact information. We excluded all the networks that previously opted out from similar measurements and did not receive any new requests. We additionally randomized our input lists across the IP space and TLDs so that no single entity receives a bulk of DNS requests at once. Our scanning activities were spread across several days.

More importantly, we reproduced previous studies at a large scale and uncovered significant security issues with tested domain names and recursive resolvers. Apart from estimating the scale of the problem, we provided recommendations that will hopefully help various stakeholders secure their systems. Therefore, we believe our large-scale measurements to have provided a benefit to the community.

6. Related Work

Researchers previously showed that a small number of domain names are cryptographically signed. Van Adrichem et al. [72] found that 7.93% from the sample of 282,766 domains under 4 TLDs deployed DNSSEC. Two years later, Wander [76] enumerated 6.4 million signed second-level domains across all the TLDs. Both verified that the great majority of domains deploying DNSSEC were signed correctly. Chung et al. [8] further analyzed 150 million domains under .com, .net, and .org, where roughly 1% were signed. In this paper, we measured more than 251 million second-level domains and found that 15.2 million attempted DNSSEC signing, more than 9 million doing it correctly.

Open DNS resolvers also received substantial attention from the research community. In 2015, Kührer et al. [46] enumerated more than 26 million open IPv4 resolvers but the collective remediation efforts decreased this number to several million by 2021-2022 [43], [44], [56], [58], [60], [78]. Hendriks et al. [28] specifically focused on the IPv6 address space and discovered 1,038 IPv6 resolvers by traversal from IPv4-only to IPv6-only zones.

Finally, existing work measured the adoption of SPF and DMARC, which appeared to be insufficient [22], [29]. Durumeric et al. [16] showed that only 35% of email servers associated with Alexa top 1 million domain names properly deploy email security mechanisms. A recent study explored the degree of SPF and DMARC deployment for high-profile domains, including banking domains, and identified misconfigured ones [49]. The authors notified domain owners through local, and national CERT teams, and as many as 23.2% of the domains were reconfigured. While it was a one-time notification campaign, such ongoing efforts to measure the deployment and raise awareness of the problem should be promoted by governments and national CERTs.

7. Conclusions

In this paper, we have evaluated the level of DNS infrastructure vulnerabilities across millions of recursive resolvers and authoritative nameservers. We have shown that they do not adequately protect against security threats such as DNS cache poisoning, email spoofing, and misuse in reflection and amplification DDoS attacks. We have enumerated the vulnerable systems and proposed a set of recommendations to registrars, registries, Internet Service Providers, and national CERT teams. They will be considered for inclusion in the upcoming legislation on EU cybersecurity. Despite our study focusing on the European Union, the proposed recommendations can also be adopted by the broader audience.

Acknowledgements

We thank Sourena Maroofi (KOR Labs) for his valuable feedback. This work has been carried out in the framework of the VIGIE 2020/0653 project funded by the European Commission. It was supported by the Grenoble Alpes Cybersecurity Institute (under contract ANR-15-IDEX-02), the French Ministry of Research projects PERSYVAL-Lab under contract ANR-11-LABX-0025-01 and DiNS under contract ANR-19-CE25-0009-01.

References

- [1] Joe Abley, Ólafur Guðmundsson, Marek Majkowski, and Evan Hunt. Providing Minimal-Sized Responses to DNS Queries That Have QTYPE=ANY. RFC 8482, 2019.
- [2] Yehuda Afek, Anat Bremner-Barr, and Lior Shafir. NXNSAttack: Recursive DNS Inefficiencies and Vulnerabilities. In *USENIX Security*, 2020.
- [3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. DNS Security Introduction and Requirements. RFC 4033, 2005.
- [4] Hadi Asghari. pyasn, 2023. <https://github.com/hadiasghari/pyasn>.
- [5] Derek Atkins and Rob Austein. Threat Analysis of the Domain Name System (DNS). RFC 3833, 2004.
- [6] Elaine Barker and Allen Roginsky. Transitions: Recommendation for Transitioning the Use of Cryptographic Algorithms and Key Lengths. In *Special Publication (NIST SP), National Institute of Standards and Technology, Gaithersburg, MD*, 01 2011.
- [7] CAIDA. AS Rank: A ranking of the largest Autonomous Systems (AS) in the Internet, 2023. <https://asrank.caida.org>.
- [8] Taejoong Chung, Roland van Rijswijk-Deij, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. A Longitudinal, End-to-End View of the DNSSEC Ecosystem. In *USENIX Security*, 2017.
- [9] Taejoong Chung, Roland van Rijswijk-Deij, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, and Christo Wilson. Understanding the Role of Registrars in DNSSEC Deployment. In *IMC*, 2017.
- [10] Cisco. OpenDNS, 2023. <https://www.opendns.com/setupguide/>.
- [11] Cloudflare. What is 1.1.1.1?, 2023. <https://www.cloudflare.com/learning/dns/what-is-1.1.1.1/>.
- [12] European Commission, Content Directorate-General for Communications Networks, Technology, J Bayer, Y Nosyk, O Hureau, S Fernandez, S Paulovics, A Duda, and M Korczyński. *Study on Domain Name System (DNS) abuse : technical report. Appendix I*. Publications Office of the European Union, 2022.
- [13] D Crocker, T Hansen, and M Kucherawy. DomainKeys Identified Mail (DKIM) Signatures. RFC 6376, 2011.
- [14] Cybersecurity & Infrastructure Security Agency. UDP-Based Amplification Attacks, 12 2019. <https://www.cisa.gov/news-events/alerts/2014/01/17/udp-based-amplification-attacks>.
- [15] D Dittrich and E Kenneally. The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. https://catalog.caida.org/paper/2012_menlo_report_actual_formatted, 2012.
- [16] Zakir Durumeric, David Adrian, Ariana Mirian, James Kasten, Elie Bursztein, Nicolas Lidzboriski, Kurt Thomas, Vijay Eranti, Michael Bailey, and J. Alex Halderman. Neither Snow Nor Rain Nor MITM...: An Empirical Analysis of Email Delivery Security. In *IMC*, 2015.
- [17] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. ZMap: Fast Internet-Wide Scanning and Its Security Applications. In *USENIX Security*, 2013.
- [18] European Commission. The EU’s Cybersecurity Strategy in the Digital Decade. <https://digital-strategy.ec.europa.eu/en/library/eu-cybersecurity-strategy-digital-decade>, 2020.
- [19] Federal Bureau of Investigation. 2020 Internet Crime Report, 2020. https://www.ic3.gov/Media/PDF/AnnualReport/2020_IC3Report.pdf.
- [20] Federal Office for Information Security. Reports on Openly Accessible Server Services, 2023. https://www.bsi.bund.de/EN/Themen/Unternehmen-und-Organisationen/Cyber-Sicherheitslage/Reaktion/CERT-Bund/CERT-Bund-Reports/Offene-Server-Dienste/offene-server-dienste_node.html.
- [21] Ondřej Filip. DNSSEC.CZ, 10 2012. <https://archive.icann.org/en/meetings/toronto2012/bitcache/DNSSEC.CZ-vid=41901&disposition=attachment&op=download.pdf>.
- [22] Ian D. Foster, Jon Larson, Max Masich, Alex C. Snoeren, Stefan Savage, and Kirill Levchenko. Security by Any Other Name: On the Effectiveness of Provider Based Email Security. In *CCS*, 2015.
- [23] Oliver Gasser, Quirin Scheitle, Paweł Foremski, Qasim Lone, Maciej Korczyński, Stephen D. Strowes, Luuk Hendriks, and Georg Carle. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *IMC*, 2018.
- [24] Google Developers. Frequently Asked Questions, 2023. <https://developers.google.com/speed/public-dns/faq>.
- [25] Google Developers. Public DNS, 2023. <https://developers.google.com/speed/public-dns>.
- [26] Wesley Griffin and Jakob Schlyter. Using DNS to Securely Publish Secure Shell (SSH) Key Fingerprints. RFC 4255, 2006.
- [27] Phillip Hallam-Baker, Rob Stradling, and Jacob Hoffman-Andrews. DNS Certification Authority Authorization (CAA) Resource Record. RFC 8659, 2019.
- [28] Luuk Hendriks, Ricardo de Oliveira Schmidt, Roland van Rijswijk-Deij, and Aiko Pras. On the Potential of IPv6 Open Resolvers for DDoS Attacks. In *PAM*, 2017.
- [29] Hang Hu, Peng Peng, and Gang Wang. Towards Understanding the Adoption of Anti-Spoofing Protocols in Email Systems. In *SecDev*, 2018.
- [30] IANA. Root KSK Ceremony 1, 06 2010. <https://www.iana.org/dnssec/ceremonies/1>.
- [31] ICANN. New gTLD Program Safeguards Against DNS Abuse, 07 2016. <https://newgtlds.icann.org/en/reviews/dns-abuse/safeguards-against-dns-abuse-18jul16-en.pdf>.
- [32] ICANN. Domain Name System Security Extensions Now Deployed in all Generic Top-Level Domains, 12 2020. <https://www.icann.org/en/announcements/details/domain-name-system-security-extensions-now-deployed-in-all-generic-top-level-domains-23-12-2020-en>.
- [33] ICANN. Knowledge-Sharing and Instantiating Norms for DNS and Naming Security. <https://kindns.org>, 2022.
- [34] ICANN. Centralized Zone Data Service, 2023. <https://czds.icann.org/home>.
- [35] ICANN Research. TLD DNSSEC Report (2021-05-18 00:05:03), 05 2021. http://stats.research.icann.org/dns/tld_report/archive/20210518.000101.html.
- [36] Internet Stiftelsen. Access to zonefiles for .se and .nu, 2023. <https://internetstiftelsen.se/en/domains/tech-tools/access-to-zonefiles-for-se-and-nu/>.
- [37] Internet stiftelsen. Recommendations for DNSSEC deployment, 2023. <https://internetstiftelsen.se/en/domains/tech-tools/recommendations-for-dnssec-deployment/>.
- [38] Internet Systems Consortium. BIND 9, 2023. <https://www.isc.org/bind/>.
- [39] Liz Izhikevich, Gautam Akiwate, Briana Berger, Spencer Drakon-taidis, Anna Ascherman, Paul Pearce, David Adrian, and Zakir Durumeric. ZDNS: A Fast DNS Toolkit for Internet Measurement. In *IMC*, 2022.
- [40] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. Millions of Targets under Attack: A Macroscopic Characterization of the DoS Ecosystem. In *IMC*, 2017.
- [41] S Kitterman. Sender Policy Framework (SPF) for Authorizing Use of Domains in Email. RFC 7208, 2014.
- [42] Olaf Kolkman, Matthijs Mekking, and R. (Miek) Gieben. DNSSEC Operational Practices, Version 2. RFC 6781, 2012.
- [43] Maciej Korczyński, Yevheniya Nosyk, Qasim Lone, Marcin Skwarek, Baptiste Jonglez, and Andrzej Duda. Don’t Forget to Lock the Front Door! Inferring the Deployment of Source Address Validation of Inbound Traffic. In *PAM*, 2020.
- [44] Maciej Korczyński, Yevheniya Nosyk, Qasim Lone, Marcin Skwarek, Baptiste Jonglez, and Andrzej Duda. Inferring the Deployment of Inbound Source Address Validation Using DNS Resolvers. In *ANRW*, 2020.
- [45] E Kucherawy, M Zwicky, and E Zwicky. Domain-Based Message Authentication, Reporting, and Conformance (DMARC). RFC 7489, 2015.

- [46] Marc Kühner, Thomas Hupperich, Jonas Bushart, Christian Rossow, and Thorsten Holz. Going Wild: Large-Scale Classification of Open DNS Resolvers. In *IMC*, 2015.
- [47] Marc Kühner, Thomas Hupperich, Christian Rossow, and Thorsten Holz. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *USENIX Security*, 2014.
- [48] Sourena Maroofi, Maciej Korczyński, and Andrzej Duda. From Defensive Registration to Subdomain Protection: Evaluation of Email Anti-Spoofing Schemes for High-Profile Domains. In *TMA*, 2020.
- [49] Sourena Maroofi, Maciej Korczyński, Arnold Hölzel, and Andrzej Duda. Adoption of Email Anti-Spoofing Schemes: A Large Scale Analysis. *IEEE Transactions on Network and Service Management*, 18(3):3184–3196, 2021.
- [50] Jared Mauch. Re: Spoofing ASNs (Re: SNMP DDoS: the vulnerability you might not know you have), 08 2013. <http://seclists.org/nanog/2013/Aug/132>.
- [51] MaxMind Developers. GeoLite2 Free Geolocation Data, 2023. <https://dev.maxmind.com/geoip/geoLite2-free-geolocation-data>.
- [52] Ministry of Industry and Trade. Digital Czech Republic v. 2.0 - The Way to the Digital Economy, 04 2014. <https://www.mpo.cz/dokument149132.html>.
- [53] Ministry of Industry, Business and Financial Affairs. Bekendtgørelse om internetdomænet .dk, January 2020. <https://www.retsinformation.dk/eli/ta/2020/44>.
- [54] P. Mockapetris. Domain names - implementation and specification. RFC 1035, 1987.
- [55] Marcin Nawrocki, Mattijs Jonker, Thomas C. Schmidt, and Matthias Wählisch. The Far Side of DNS Amplification: Tracing the DDoS Attack Ecosystem from the Internet Core. In *IMC*, 2021.
- [56] Marcin Nawrocki, Maynard Koch, Thomas C. Schmidt, and Matthias Wählisch. Transparent Forwarders: An Unnoticed Component of the Open DNS Infrastructure. In *CoNEXT*, 2021.
- [57] Yevheniya Nosyk, Maciej Korczyński, and Andrzej Duda. Routing Loops as Mega Amplifiers for DNS-Based DDoS Attacks. In *PAM*, 2022.
- [58] Yevheniya Nosyk, Maciej Korczyński, Qasim Lone, Marcin Skwarek, Baptiste Jonglez, and Andrzej Duda. The Closed Resolver Project: Measuring the Deployment of Source Address Validation of Inbound Traffic. *IEEE/ACM Transactions on Networking*, 2023.
- [59] OVH. Securing your domain name with DNSSEC, 10 2022. https://docs.ovh.com/gb/en/domains/secure_your_domain_with_dnssec/.
- [60] Jeman Park, Rhongho Jang, Manar Mohaisen, and David Mohaisen. A Large-Scale Behavioral Analysis of the Open DNS Resolvers on the Internet. *IEEE/ACM Transactions on Networking*, 30(1):76–89, 2022.
- [61] Craig Partridge and Mark Allman. Ethical Considerations in Network Measurement Papers. *Commun. ACM*, 59(10):58–64, sep 2016.
- [62] PeeringDB. The Interconnection Database, 2023. <https://www.peeringdb.com>.
- [63] Quad9. A public and free DNS service for a better security and privacy, 2023. <https://www.quad9.net>.
- [64] Scott Rose, Matt Larson, Dan Massey, Rob Austein, and Roy Arends. Protocol Modifications for the DNS Security Extensions. RFC 4035, 2005.
- [65] Scott Rose, Matt Larson, Dan Massey, Rob Austein, and Roy Arends. Resource Records for the DNS Security Extensions. RFC 4034, 2005.
- [66] Christian Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *NDSS*, 2014.
- [67] SIDN. DNSSEC - A cryptographic security extension to the DNS protocol, October 2021. <https://www.sidn.nl/en/faq/dnssec>.
- [68] Marcin Skwarek, Maciej Korczyński, Wojciech Mazurczyk, and Andrzej Duda. Characterizing Vulnerability of DNS AXFR Transfers with Global-Scale Scanning. In *IEEE Security and Privacy Workshops (SPW)*, 2019.
- [69] Raffaele Sommese, Mattijs Jonker, and KC Claffy. Observable KINDNS: Validating DNS Hygiene. In *IMC*, 2022.
- [70] The European Parliament and The Council of The European Union. Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS 2 Directive) (Text with EEA relevance). <https://eur-lex.europa.eu/eli/dir/2022/2555>, 2022.
- [71] University of Oregon. University of Oregon Route Views Project, 2023. <http://www.routeviews.org/routeviews/>.
- [72] Niels L. M. van Adrichem, Norbert Blenn, Antonio Reyes Lua, Xin Wang, Muhammad Wasif, Ficky Fatturrahman, and Fernando A. Kuipers. A Measurement Study of DNSSEC Misconfigurations. *Security Informatics*, 4:1–14, 2015.
- [73] Olivier van der Toorn, Johannes Krupp, Mattijs Jonker, Roland van Rijswijk-Deij, Christian Rossow, and Anna Sperotto. ANYway: Measuring the Amplification DDoS Potential of Domains. In *CNSM*, 2021.
- [74] Roland van Rijswijk-Deij, Anna Sperotto, and Aiko Pras. DNSSEC and Its Potential for DDoS Attacks: A Comprehensive Measurement Study. In *IMC*, 2014.
- [75] Saurabh Verma, Ali Hamieh, Jun Ho Huh, Henrik Holm, Siva Raj Rajagopalan, Maciej Korczyński, and Nina Fefferman. Stopping Amplified DNS DDoS Attacks through Distributed Query Rate Sharing. In *ARES*, 2016.
- [76] Matthias Wander. Measurement Survey of Server-Side DNSSEC Adoption. In *TMA*, 2017.
- [77] Paul Wouters and Ondřej Surý. Algorithm Implementation Requirements and Usage Guidance for DNSSEC. RFC 8624, 2019.
- [78] Ramin Yazdani, Roland van Rijswijk-Deij, Mattijs Jonker, and Anna Sperotto. A Matter of Degree: Characterizing the Amplification Power of Open DNS Resolvers. In *PAM*, 2022.