

Characterizing Vulnerability of DNS AXFR Transfers with Global-Scale Scanning

Marcin Skwarek*, Maciej Korczyński[†], Wojciech Mazurczyk*, and Andrzej Duda[†]

*Warsaw University of Technology, 00-665 Warsaw, Poland

Email: mskwarek@stud.elka.pw.edu.pl, wmazurczyk@tele.pw.edu.pl

[†]Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, F-38000 Grenoble, France

Email: firstname.lastname@univ-grenoble-alpes.fr

Abstract—In this paper, we consider security issues related to zone transfers by investigating the responses of DNS servers to AXFR requests. In particular, we investigate how attackers can exploit available AXFR zone transfers to obtain useful reconnaissance data. To evaluate the extent of the security flaw, we have scanned DNS servers on a global scale with a dedicated tool and transferred multi-line zone files of 3.6M domains. We have first analyzed the experimental data to evaluate the size of the DNS zones. Then, we have investigated what kind of information zone transfers may reveal to attackers. We have also studied the information on chosen services that attackers can use in further attacks and analyzed potential security problems such as enumerating open SMTP relays or domains vulnerable to DNS hijacking. Finally, we have proposed potential remediation strategies to improve the security of the DNS ecosystem.

Index Terms—AXFR Protocol, Zone File Transfer, Domain Name System, Internet Measurements, Network Security

I. INTRODUCTION

Reconnaissance is a type of network attacks with the main aim to gather as much information on a targeted victim as possible to identify potential vulnerabilities. In *passive reconnaissance* techniques, an attacker only monitors network traffic without interfering with target devices. *Active reconnaissance* techniques involve malicious actions that interact with the victim system, for example, by transmitting intentionally crafted packets or by performing port scanning.

In this paper, we focus on an active reconnaissance technique relying on DNS and its Asynchronous Transfer Full Range (AXFR) mechanism for zone transfers [23]. A zone transfer is the main DNS component for replication of DNS data across a set of DNS servers. It relies on TCP and a client-server communication model. In a common setup, a client is a slave DNS server requesting DNS data from a master DNS server. The server sends a zone file, i.e., the DNS information represented as a text file in response to the client request.

The DNS zone information may include sensitive information about the internal infrastructure of a given system or network useful for an attacker for performing direct attacks such as Distributed Reflection Denial of Service (DRDoS) [27] or DNS zone poisoning [20]. The attacker can, for example, enumerate hosts running outdated and potentially vulnerable operating systems (by analyzing HINFO resource records), open mail servers that can be used as SMTP relays to distribute spam (by analyzing MX and TXT resource records), or DNS,

NTP or SMTP servers that can be used as amplifiers in DRDoS attacks. For example, after finding potentially misconfigured NTP servers (based on the content of transferred zone files), an attacker can check if they support the `monlist` request. Such misconfigured servers share the list of their recent clients in up to 100 UDP datagrams with 440 byte payload each, resulting in the DRDoS bandwidth amplifier factor equal to 556 on the average [27].

To avoid such security leaks, all authoritative name servers have to restrict zone transfers to authorized DNS servers only. The most secure configuration is to allow zone transfers to authenticated authoritative servers using Secret Key Transaction Authentication (TSIG) [32].

Note that some domain owners or organizations may decide to make their zone files visible to the Internet public for transparency reasons. The Internet Foundation in Sweden—the registry operator of `.se` and `.nu` country-code top-level domain names (ccTLDs)—makes its zone files accessible for all Internet users [17]. However, zone accessibility is limited to the information about the 2nd-level domain delegations including DNSSEC-related records and does not provide information about the network and DNS infrastructure of subdomains.

CERT advisories and blogs have already identified the problem of unrestricted DNS zone transfers as a vulnerability (CVE-1999-0532), but its relevance in the global DNS landscape has not been studied yet [12], [14], [29], [34].

In this paper, we consider the security issues related to zone transfers by investigating the responses of DNS servers to AXFR transfer requests. In the global-scale experiments, we have transferred 6.1M multi-line zone files and revealed 62M unique DNS records that correspond to 3,6M domains. We have also studied how cybercriminals can benefit from the information gathered in similar scans and propose potential remediation strategies. Our research goal is to strengthen the security of the DNS ecosystem via the results of our experiments and subsequent notifications to affected parties. We make the source code of our scanner available at <https://github.com/mskwarek/myDig> to encourage reproducibility.

II. RELATED WORK

Previous work has focused on establishing whether authoritative DNS servers for the Alexa Top 1M domains [1] are vulnerable to zone transfers [4]. The main difference between

the work presented in this paper and the existing research is the scale of the performed experiment. The evaluation of only popular domains is a serious limitation that results in a partial view of the DNS landscape.

A popular study of the security issues related to zone transfers concerned a misconfigured name server at Western Digital that allowed unauthorized DNS zone transfers [9]. As a consequence, 1.1M addresses of clients that used the *My Cloud NAS* product leaked. In our work, we focus on a significant number of name servers on the Internet and we do not limit our experiments to a single company or a single product.

Another vulnerability related to our work is so-called *zone enumeration*. Initially, ability to enumerate all the names in a zone was not considered as an error [2], however, later on, a compromise has been reached [22] that in certain cases, the knowledge of all the domain names in a zone can lead to security risks. RFC 5155 [22] provides some examples showing that due to this vulnerability, it is easier for an attacker to obtain email addresses for future spam campaigns or data useful during a reconnaissance phase of a network attack.

III. EXPERIMENTAL METHODOLOGY

A. Bootstrapping DNS Data

To request an AXFR transfer, a client needs to know the domain name and its authoritative name server. We have first leveraged A and NS resource records (RRs) for all domains observed in three complementary datasets: DNSDB—a large database of passively observed DNS queries fed by hundreds of nodes across the world generously provided by Farsight Security [11], DNS ANY responses for known forward DNS names stored in the Internet-Wide Scan Data Repository hosted by Censys [5], and available zone files. We have obtained the .com, .net and .name zone files from Verisign [31], performed zone transfers to replicate DNS databases of the .se and .nu ccTLD [17] and .nl zone file (under the contract of SIDN—the .nl ccTLD registry). We have also collected zone files from the .us ccTLD, .biz, .org, .asia, .info, .mobi, .post, and .tel legacy gTLDs and 1,230 new gTLDs made available through the Centralized Zone Data Service by ICANN [15]. Finally, we have added the dataset with the domains listed in the Alexa Top 1M Global Sites [1]. Although previous works have shown that Alexa is vulnerable to large-scale manipulation [26], [28], it is still one of the most widely used metrics that characterizes the website’s popularity. Another limitation is that DNSDB may contain poisoned [7] obsolete or incorrect records, but we aimed at creating a possibly largest and a most complete overview of the global domain space.

We have extracted 2nd-level domain names (and upper-level domains if a given registry provides such registrations, e.g., example.co.uk [21]), their name servers, and the IP addresses of name servers. We then performed active DNS queries to obtain missing data if for a given domain, the authoritative server was not passively observed in DNSDB or if in the zone files, DNS glue records were missing. We have then excluded invalid domains or the domains that resolved to the special ICANN IP addresses 127.0.53.53 indicating that a

name collision occurred [16], all .arpa domains, the domains resolving to IP addresses of private networks or invalid IP addresses, and the domains/IP addresses of networks managed by administrators that contacted us in the past asking to exclude them from Internet-wide measurements. For the total 353,870,510 unique domains in the aggregated datasets, we have enumerated all combinations of the corresponding name servers and their IP addresses (3,855,615 in total), and finally created a list of 5,032,117,394 *domain, name server IP address* pairs used in our measurement campaign.

B. Scanner Dedicated to AXFR Transfers

We have developed a scanner for requesting zone transfers at the scale of the Internet. Its goal was to operate in a highly efficient way compared to existing standard tools such as *dig*. The scanner core handles command line parsing, reading and writing DNS packets, and parsing the text files containing domain and name server pairs to scan. The scanner generates the application protocol data unit that complies with the standard DNS message format defined in RFC 1035 [25]. It supports both UDP and TCP transport-layer protocols and the most frequently used DNS query types. It was written in C and tested on GNU/Linux.

C. Ethical Considerations

We have submitted a Research Ethics Application to the Human Research Ethics Committee. The committee approved our request under condition that we will not publish the collected data and apply the principles to ICT Research as described in (the Companion to) the Menlo Report [8].

We have promoted full transparency in the study, its objectives, and a clear opt-out mechanism on the site hosted on our measurement server. All questions could be sent to our email address provided on our website. During our measurement campaigns, 7 organizations contacted us to include their resources on the do-not-scan list.

IV. EXPERIMENTAL RESULTS

We performed the global-scale vulnerability assessment for 5,032,117,394 *domain, name server IP address* pairs in Jan 2018 and successfully transferred 11,366,058 DNS zones. As expected, many servers did not respond. In addition to some obsolete NS information, this effect can also indicate network filtering. The results give us thus a lower bound on vulnerability to AXFR transfers in the global domain name space.

A. Types of Responses

We have observed a variety of responses to AXFR requests which can be divided into one-line (i.e. with only on RRs) or multi-line zone files (more than one RRs). The entry present in one-line zones is of the SOA type. We have downloaded 5,232,253 single- and 6,133,805 multi-line zone files. The multi-line zones are more interesting from the attacker’s point of view as they may reveal more meaningful information. We therefore focus below on the analysis of multi-line zone files.

In multi-line zones, we enumerated in total 61,955,666 unique RR of 59 different types (A, NS, MX, etc.). Note that




Figure 1. Cumulative % of collected resource records as a function of the number of RRs in multi-line zone files ordered by their size.

90% of multi-line zone files had less than 26 RRs each, with the mean and median sizes equal to 17 and 13, respectively. Such zones are typical for domains used by individual users or small organizations. In their zones, we usually find one IP address, under which a server provides basic services such as a web or mail server, as well as the entries for DNS servers.

Fig. 1 presents a cumulative % of all collected RRs as a function of the number of RRs in multi-line DNS zones. Note that a relatively smaller number of larger zone files (represented by a long tail) reveals a significant number of RRs. We have transferred 46,293 zones larger than 100 entries and 3,279 zones with the size larger than 1000 entries. The largest zone, maintained by a hosting company, contained 2,698 RRs. At the time of writing, some zones have expanded significantly, with the largest one containing 18,922 RRs.

Large zone files usually contain more diversified entries such as SPF (Sender Policy Framework) encapsulated in the TXT records, or DNSKEY and RRSIG entries, indicating that a given zone uses DNSSEC.

B. Affected Domains

To get insight into the type of vulnerable domains, we first compare the proportion of all affected domains with the vulnerable Alexa Top 1M popular websites. Among all 353,870,440 scanned domains, we successfully transferred zones of 3,604,371 domains (1.02%). Surprisingly, we find a higher proportion (2.77% corresponding to 27,736 domains) for the Alexa Top 1M. However, we should note that 1.02% establishes a conservative lower bound for the magnitude of the problem as many domains might have expired before our measurement study. The most popular affected domains reached 4th and 6th (baidu.com and qq.com) ranks on the Alexa list. At the time of writing, however, in both cases their respective authoritative name servers do not support unrestricted AXFR transfers anymore. It implies that those were misconfigured rather than configured that way on purpose.

C. Affected Name Servers

We have further investigated if master, slave or both types of authoritative name servers are vulnerable to AXFR transfers. We first collected SOA records for all domains that returned multi-line zones to determine their master servers. In a regular configuration, only the master server should be configured to permit AXFR transfers to authorized clients. Surprisingly, we performed zone transfers using only masters for 553,811 domains (15.4%) and we transferred zone files for 1,088,219

domains (30.2%) using both master and slave servers. Finally, in 1,962,341 (54.4%) domains, only slave servers enabled unrestricted zone transfers. The high number of slaves responding to AXFR requests indicates that unrestricted transfers are the result of misconfiguration rather than purposeful action aimed at making the DNS information public.

D. Revealed Services

We have further analyzed how much information can be obtained on the services running on the machines found in the downloaded zone files. This type of information allows an attacker to gain knowledge about the infrastructure and enables preparation of more complex network attacks. During the experiments, we have leveraged the information about the following services: mail servers, file transfer (FTP), network time protocol (NTP), version-control systems (VCS): git and svn, continuous integration systems (CI): jenkins and gitlab, test versions of web services, DNS servers, operating systems (OSs), and IPv6-enabled hosts.

We have extracted the information about the services by analyzing collected fully qualified domains names (FQDN) in which the lowest-level domain suffix contained appropriate keywords, i.e., *mail*, *ftp*, *ntp*, *git*, *svn*, *jenkins*, *gitlab*, *dev*, and *test* (e.g., a domain ftp.domain.com is assumed to provide the FTP service). The exceptions are DNS, mail servers, OSs, and IPv6 hosts for which the relevant information was also gathered from the NS, MX, HINFO, and AAAA RRs. The proposed heuristics based on matching the lowest-level domain suffixes with a predefined list of keywords do not allow to identify all individual systems. However, it has been possible to retrieve important information about the running services and to assess to which extent it may be used for nefarious purposes.

1) *Mail Servers*: We have identified 2,578,948 domains with the lowest-level domain suffix being *mail*. In the zone files, we also enumerated 4,563,557 unique MX RRs containing FQDNs of the mail servers. We hypothesize that a set of mail servers obtained through the zone transfers may be less secure because their administrators allow zone transfers, which is considered a bad practice.

An attacker can use the information about mail servers in several ways. She may try to verify whether valid email accounts exist on a server by connecting to the server on port 25 and running the SMTP VRFY command [19]. Spammers can automate this method to perform a directory harvest attack (DHA)—a way of generating valid mail addresses from a given domain name using brute force. Attackers may also check if the mail servers are open SMTP relays and can use them to send outbound spam emails, or generate DRDoS attacks [27]. To empirically demonstrate one of the threats, we randomly selected 20,921 SMTP servers, and using `nmap`¹ we enumerated 475 open SMTP relays (2.3% of all scanned servers) and notified their operators about the vulnerability.

2) *Sender Policy Framework*: SPF is an email authentication method designed to detect spoofed sender addresses [18].

¹<https://nmap.org/nsedoc/scripts/smtp-open-relay.html>

SPF information is encapsulated in the DNS TXT records and allows to define rules based on which a mail server accepts an incoming email or not. Qualifiers are important SPF elements. In particular, the `v=spf1 +all` qualifier indicates that for a given domain, the incoming email should be always accepted.

We have discovered 1.8M records with SPF information and 742 of them had the least restrictive `v=spf1 +all` qualifier activated, which provides important information to the organizers of spam or phishing campaigns. If a malicious party knows that for the certain domain, SPF rules will not filter out any emails, then an attacker can send mass mailing to all its victims and be sure that spam messages will reach intended recipients thus saving time (assuming that no other security measures are used).

To assess the vulnerability, we have configured our own mail server and attempted to send emails to our test Yahoo mailbox (that authenticates emails with SPF) with spoofed sender addresses. For example, if the TXT RR enumerated in the zone is `example.com 38400 16 v=spf1 +all`, then we send a test email from `test@example.com`. We have randomly selected 49 out of 742 vulnerable domains for which all of our test emails were successfully delivered. When we have sent mails using domains with more restrictive SPF qualifiers, the mail delivery failed because our mail server was not authorized as a permitted sender.

3) *FTP servers*: FTP was the second most popular service (after mail servers) in the collected zone data—we found 2,027,651 unique instances of FTP servers. The attacker may try to connect anonymously if the server's configuration allows it or to enumerate FTP banners for finding important information such as the software version and then try to identify its vulnerabilities against known exploits. To demonstrate the issue, we have randomly sampled 3,468 FTP servers, and enumerated (using `nmap`) 33 allowing anonymous FTP login (0.95% of all scanned servers). We informed their respective administrators about the problem.

In early 2018, one of the email campaigns abused hacked FTP servers as download locations for Dridex Trojan—one of the most prolific banking malware in recent years [3].

4) *NTP servers*: We have also identified 2,619 NTP servers. Attackers can, for example, perform an additional smaller-scale, less conspicuous vulnerability scan to verify if they are publicly accessible and if they respond to the NTP Mode 6 or 7 queries. Such misconfigured NTP servers respond with much larger responses than queries, so they can be used in DRDoS attacks with the bandwidth amplifier factor equal to 556 [27].

5) *VCS and CI Servers*: In software engineering, VCS and CI servers are used to track and provide control over changes to the source code. Our data revealed 3,955 git servers, 1,391 gitlab environment servers, 3,630 svn, and 681 jenkins systems.

An attacker can determine if some common software development tools appear in the data from the transferred zones. When a company runs its own local version of VCS such as GitHub, GitLab or BitBucket, an attacker can steal some intellectual property of the company. GitLab, for example,

suffers from a number of vulnerabilities that can be exploited². Thus, the information from a zone transfer enables an attacker to list potentially unpatched systems.

CI servers, such as Jenkins, are another target of attacks. The most burdensome attack vector would be to launch a DDoS attack against a CI server because it may disrupt the workflow of a company. The list of Jenkins software vulnerabilities is also long and many of the identified flaws do not even require for an attacker to log into the targeted system³.

6) *Test Domains*: Another aspect we investigated was the occurrence frequency of test and developers domains. In the collected data, we have discovered 36,997 domain names with *dev.* and 48,086 domains with *test.* prefixes.

Such domains are usually not properly secured as they serve to evaluate, for example, security mechanisms. Compromising these type of domains can further enable access to other machines within the targeted network, which offers a variety of possibilities to mount successful attacks.

7) *DNS Servers*: We have also investigated establishing a list of other DNS servers. By analyzing the information collected in the transferred zones, we can determine RRs that delegate resolving domain names to other name servers.

Based on the collected measurement data, we have established a list of 480,504 unique, previously unseen name servers (and domain names for which they are authoritative for) that an attacker can scan to determine vulnerable ones, i.e., those that have commonly known security flaws.

To empirically demonstrate how the obtained NS records can be used by the attacker, we identified domains vulnerable to zone poisoning [20]. This critical vulnerability allows anyone who can reach an authoritative name server of a given domain to update its zone file and, for example, hijack the domain name. We followed the measurement methodology and ethical principles as outlined by Korczyński et al. [20]. Our scanner attempts to add an extra A RR to the zone file, associating a new upper-level suffix, i.e., *research*, with the IP address of the web server of our project. For example, if the NS record enumerated in the transferred zone is `example.com NS ns1.example.com`, then we try to add an extra A record for a new subdomain, i.e., `research.example.com A 1.2.3.4`.

From the initial set of 480,504 domains, only 190,493 were still registered. Using public Google, Dyn, and Quad9 DNS resolvers, we have enumerated the IP addresses of the authoritative name servers. For each domain, we sent an UPDATE request directly to all IP addresses of name servers. Then, we sent an A request to each of the potentially updated servers to verify if the zone was indeed updated. We observed 283 successfully added A RRs, corresponding to 80 unique name servers and 219 unique domain names (0.046% of newly discovered and active domains). All added records were successfully deleted after the study. We also notified all operators of the servers vulnerable to zone poisoning.

² www.cvedetails.com/vulnerability-list/vendor_id-13074/Gitlab.html

³ www.cvedetails.com/vulnerability-list/vendor_id-15865/Jenkins.html

The information obtained by the attacker through AXFR transfers may lead to complex cyberattacks such as the recently discovered global DNS infrastructure hijacking campaign [30].




Figure 2. Distribution of OSs gathered from HINFO RRs of transferred zones.

E. Host Information

The HINFO RR typically stores the information about the host operating system and CPU details. In the collected measurements data, we have found 9,480 unique HINFO records. In some cases, the investigated machines ran obsolete versions of Windows OS (see Fig. 2). Among 1,291 revealed Windows machines, we found 328 Windows 98, 328 Windows XP, and 235 Windows NT. They raise a serious security risk as those OS versions are no longer supported by Microsoft, making them easy targets for attackers.

From these records we have also obtained information about the network infrastructure and revealed, for example, devices running Cisco Aironet 1200, Cisco 3725, Cisco Catalyst 3750G-24TS-S, or Catalyst 2960PD-8TT-L.

F. IPv6 Addresses

Due to the size of the IPv6 address space, exhaustive scans of the entire space are not possible so it is more difficult to find the addresses of active hosts. Therefore, an attacker can be interested in generating target lists of IPv6 addresses (called IPv6 hitlists [13]). She can further scan the discovered machines to identify security issues. Our measurements revealed 290,300 unique IPv6 addresses in AAAA records in the collected zone files.

V. REMEDIATION STRATEGIES

As zone transfers expose information enabling further attacks, we need to find effective mechanisms to inform the affected parties. After the introduction of the General Data Protection Regulation (GDPR), the Registrant and Administrative Contact is no longer displayed in the public WHOIS. Therefore, we cannot directly contact the owners of affected domains. An alternative method is to reach out to intermediaries such as national CERTs or registry operators of TLDs. In this subsection, we investigate the distribution of vulnerable zones over Autonomous System (AS) operators, countries, and TLDs and try to propose the most effective way to reach the parties affected by the AXFR vulnerability at scale.

Fig. 3 shows the distribution of name servers over AS operators ordered by the number of transferred zones (from the lowest to the highest). We have mapped all IP addresses




Figure 3. Cumulative % of multi-line zones distributed over ASes ordered by the number of transferred zones.

of name servers to their respective ASs using PyASN [10] and found that for almost 50% of transferred zones, their respective authoritative name servers are located in different ASs. The mean and median number of vulnerable name servers per AS is equal to 3 and 2, respectively. All in all, we have discovered vulnerable servers in 10,605 ASs. The widespread distribution of the name servers might be a major obstacle for remediation, as different AS operators would need to reconfigure their individual name servers to restrict AXFR transfers to the authorized clients. Therefore, a comprehensive notification campaign to AS operators might not be effective.

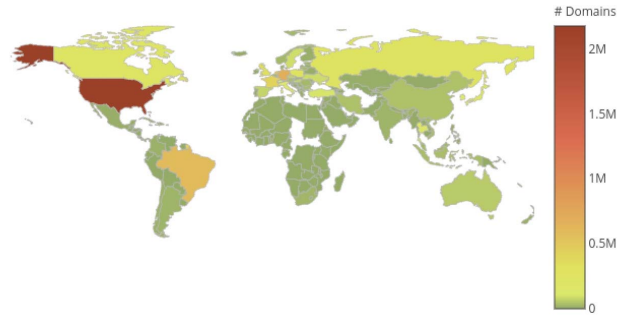


Figure 4. Distribution of vulnerable resources across different countries.

We have continued to work on this issue by investigating the distribution of vulnerable resources across different countries. Fig. 4 presents the geographic distribution of vulnerable domain, name server IP address pairs. We have used the MaxMind GeoIP database to assign IP addresses to countries [24]. We have discovered vulnerable servers in 183 countries. A relatively low number of vulnerable servers in the African region can be explained by an underdeveloped Internet infrastructure. We have found that 3,600,414 (58,58%) of vulnerable resources are located in the top 5 most affected countries (USA, Germany, Brazil, Netherlands, and France). Moreover, reaching out to only US-CERT would increase the scalability of our notifications significantly as 1,752,348 (28,51%) of all vulnerable resources are located in USA.

Finally, we have analyzed the distribution of vulnerable domains over TLDs as registries may also help in the remediation. We do not expect that they impose a policy on vulnerable domains, but rather they can notify administrators about the issue as they have access to registrar and administrative contacts no longer available publicly. Fig. 5 shows the distribution of vulnerable domains (in red) in comparison

to the TLD market share (in blue). As expected, we find that the majority (34.64%) of all vulnerable domains have a `.com` extension. Contacting only Verisign that manages a number of legacy, country-code, and new generic TLDs (`.com`, `.net`, `.tv`, `.name`, `.cc`, `.edu`, `.gov` and `.jobs`), would cover 39,17%, which means 2,407,761 domains.




Figure 5. Distribution of vulnerable domains over TLDs compared to the TLD market share [33].

By analyzing the distributions of vulnerable resources per different intermediaries, we wanted to determine the most effective way to reach out to possibly the highest number of affected parties at scale. It is not a trivial task as retrieving the contact information proved to be highly problematic [6], especially in the context of the GDPR introduction. The next step for this work is to notify affected parties via national CERTs and registry operators to improve the DNS ecosystem.

VI. CONCLUSIONS

In this paper, we have investigated how attackers can exploit available AXFR zone transfers to obtain useful reconnaissance data. To evaluate the extent of this security flaw, we have scanned DNS servers on a global scale with a dedicated tool. The results show that there are many slave and master servers allowing AXFR zone transfers, which raises important security concerns. We have collected 62M unique DNS resource records enumerated in 6.1M multi-line zones that revealed information corresponding to 3.6M domains (including 27.7K high-value domains listed in Alexa Top 1M). We have also studied the information on chosen services that attackers can use in further attacks and analyzed potential security problems. We have leveraged the information about mail, FTP, NTP, VCS, CI, and DNS servers, operating systems, and IPv6-enabled hosts. We found, for example, 4.6M MX records containing FQDNs of mail servers and through active vulnerability measurements we demonstrated that on average 2.3% represent open SMTP relays which the attacker may use to distribute spam. We enumerated 480K previously unseen NS RRs, assessed the potential impact of non-secure dynamic updates, and found 219 domains that could be hijacked using the zone poisoning attack [20]. Finally, we proposed notification strategies to improve the security of the DNS ecosystem.

The next step for this work is to start regular measurements and notify all vulnerable parties affected by AXFR transfers.

ACKNOWLEDGEMENTS

The authors would like to thank Farsight Security for providing access to the DNSDB data, SURFnet for providing access to their measurement servers and anonymous reviewers for their valuable comments.

REFERENCES

- [1] Alexa, "Actionable Analytics for the Web," <http://www.alexa.com>.
- [2] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security Introduction and Requirements," RFC 4033, 2005.
- [3] I. Arghire, "Dridex Campaign Abuses FTP Servers," www.securityweek.com/dridex-campaign-abuses-ftp-servers, 2018.
- [4] Censys Team, "Internet-Wide Scan Data Repository," <https://scans.io/study/hanno-axfr>, 2015.
- [5] —, "Internet-Wide Scan Data Repository: DNS Records (ANY)," <https://scans.io/study/sonar.fdns>, 2017.
- [6] O. Cetin, C. Ganan, M. Korczyński, and M. van Eeten, "Make Notifications Great Again: Learning How to Notify in the Age of Large-Scale Vulnerability Scanning," in *Proc. of WEIS*, 2017.
- [7] D. Dagon, N. Provos, C. P. Lee, and W. Lee, "Corrupted DNS Resolution Paths: The Rise of a Malicious Resolution Authority," in *NDSS*, 2008.
- [8] D. Dittrich and E. Kenneally, "The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research," U.S. Department of Homeland Security, Tech. Rep., August 2012.
- [9] E. Kovacs, "Western Digital User Data Exposed by DNS Issue," <https://www.securityweek.com/western-digital-user-data-exposed-dns-issue>, 2016.
- [10] EconSec group, TU Delft, "PyASN 1.6.0b1 Offline IP address to ASN lookup module," <https://pypi.python.org/pypi/pyasn>, 2017.
- [11] Farsight Security, "DNS Database (DNS-DB)," <https://www.dnsdb.info>.
- [12] G. Hill, "DNS Zone Transfer Attack," 2012, <https://security.stackexchange.com/questions/10452/dns-zone-transfer-attack>.
- [13] O. Gasser, Q. Scheitle, P. Foremski, Q. Lone, M. Korczyński, S. D. Strowes, L. Hendriks, and G. Carle, "Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists," in *Proc. of ACM IMC*, 2018.
- [14] A. Householder, B. King, K. Silva, and C. Liu, "Securing an Internet Name Server," 2002, https://resources.sei.cmu.edu/asset_files/WhitePaper/2002_019_001_52496.pdf.
- [15] ICANN, "Centralized Zone Data Service," <https://czds.icann.org>.
- [16] —, "Name Collision Occurrence Management Framework," <https://www.icann.org>, 2014.
- [17] Internet Foundation in Sweden, "Access to Zone Files for .se and .nu," <https://www.iis.se/english/domains/tech/zonefiles>.
- [18] S. Kitterman, "Sender Policy Framework (SPF) for Authorizing Use of Domains in Email," RFC 7208, 2014.
- [19] J. C. Klensin, "Simple Mail Transfer Protocol," RFC 5321, 2008.
- [20] M. Korczyński, M. Król, and M. van Eeten, "Zone Poisoning: The How and Where of Non-Secure DNS Dynamic Updates," in *Proc. of ACM IMC*, 2016, pp. 271–278.
- [21] M. Korczyński, S. Tajalizadehkhoob, A. Noroozian, M. Wullink, C. Hesselman, and M. v. Eeten, "Reputation metrics design to improve intermediary incentives for security of tlds," in *Proc. of IEEE Euro S&P*, April 2017, pp. 579–594.
- [22] B. Laurie, G. Sisson, R. Arends, and D. Blacka, "DNS Security Hashed Authenticated Denial of Existence," RFC 5155, 2008.
- [23] E. Lewis and A. Hoenes, "DNS Zone Transfer Protocol (AXFR)," RFC 5936, 2010.
- [24] MaxMind, "GeoIP Database," www.maxmind.com/en/ip-location, 2018.
- [25] P. Mockapetris, "Domain names - Implementation and Specification," STD 1035, 1987.
- [26] V. L. Pochat, T. van Goethem, S. Tajalizadehkhoob, M. Korczyński, and W. Joosen, "Tranco: A research-oriented top sites ranking hardened against manipulation," in *Proc. of NDSS*, 2019.
- [27] C. Rossow, "Amplification Hell: Revisiting Network Protocols for DDoS Abuse," in *In Proc. of NDSS*, 2014.
- [28] Q. Scheitle, O. Hohlfeld, J. Gamba, J. Jelten, T. Zimmermann, S. D. Strowes, and N. Vallina-Rodriguez, "A long way to the top: Significance, structure, and stability of internet top lists," in *Proc. of IMC*, 2018.
- [29] US-CERT, "DNS Zone Transfer AXFR Requests May Leak Domain Information," <https://www.us-cert.gov/ncas/alerts/TA15-103A>, 2015.
- [30] —, "Alert (AA19-024A): DNS Infrastructure Hijacking Campaign," <https://www.us-cert.gov/ncas/alerts/AA19-024A>, 2019.
- [31] Verisign, "TLD Zone File Access for .com, .net and/or .name," https://www.verisign.com/en_US/channel-resources/domain-registry-products/zone-file/index.xhtml.
- [32] P. Vixie, O. Gudmundsson, D. Eastlake, and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)," RFC 2845, 2000.
- [33] W3Techs, "Usage of Top Level Domains for Websites," https://w3techs.com/technologies/overview/top_level_domain/all.
- [34] R. Wood, "Zonettransfer.me," digi.ninja/projects/zonettransferme.php.