A Comprehensive, Longitudinal Study of Government DNS Deployment at Global Scale

Rebekah Houser
University of Delaware
rlhouser@udel.edu

Shuai Hao
Old Dominion University
shao@odu.edu

Chase Cotton
University of Delaware
ccotton@udel.edu

Haining Wang
Virginia Tech
hnw@vt.edu

***Abstract*—Within the Domain Name System (DNS), government domains form a particularly valuable part of the namespace, representing trusted sources of information, vital services, and gateways for government personnel to engage in their duties. As the COVID-19 pandemic has unfolded, governments' digital resources have become increasingly important to provide support to populations largely in isolation. The accessibility of these resources relies largely on the trustworthiness of the domains that represent them. In this paper, we conduct an extensive measurement study focused on the availability and legitimacy of DNS records in the authoritative nameservers of government domains for over 190 countries. Our measurements reveal that thousands of domains do not use replicated authoritative nameservers, as well as a substantial increase in the trend of more domains relying on a single third-party DNS services provider. We also find more than 1,000 domains vulnerable to hijacking due to defective delegations. Our work shows that although robust overall, the deployments of authoritative nameservers in government domains still contain a non-trivial number of configurations that do not meet RFC requirements, leading to poor performance and reduced reliability that may leave domains vulnerable to hijacking.**

***Index Terms*—E-Government, Government DNS, DNS deployment, DNS delegation**

I. Introduction

"Running a nameserver is not a trivial task" [1]. The introduction to RFC 1912 (Common DNS Operational and Configuration Errors) highlights the complexities involved in operating authoritative DNS (ADNS) nameservers. The challenge described in some ways explains the story that has played out in the DNS, where common misconfigurations persist despite many efforts to uncover their prevalence and impact and to provide solutions. Errors in the DNS are easily made, and sometimes difficult to detect and correct, as redundancy often hides the existence or impact of a single failing component. Even beyond issues of configurations that are easily recognized as incorrect (*e.g.*, delegating to authoritative servers that do not exist), DNS operators face subjective decisions about the appropriate configuration for their servers. In some scenarios, an ideal solution may not exist. In light of these realities, we have little reason for surprise that many of the errors that were featured 26 years ago in RFC 1912 continue within the DNS to this day.

While the DNS has continued to operate successfully in the presence of relatively widespread misconfigurations, the past few years have witnessed the DNS targeted or leveraged in new or increasingly severe attacks on individuals, organizations, and governments [2]–[6]. The latter point is of particular concern, as e-government plays a vital role in the daily lives of citizens. Since the COVID-19 pandemic has unfolded, information and communication technologies (ICTs) have become critically important to governments disseminating information and services as populations have grown more isolated [7]. The DNS plays a key role in this context. The government resources that citizens seek to access are represented by domain names, and the overall trustworthiness of these services necessarily rests on the reliability of authoritative DNS deployments. The ever-increasing growth in attacks against (and increasing reliance upon) the DNS calls for renewed attention to the reliability and security of the DNS, particularly those domains used by governments.

In this paper, we conduct an extensive measurement study to assess the health of the DNS with respect to domains belonging to e-government resources. We first compile a list of global government domains by extending the reported domains from the United Nations' E-Government Knowledge Base [8]. We then retrieve the relevant records from a passive DNS (PDNS) database covering a 10-year period to obtain their nameserver information and identify their evolution. In the meantime, we conduct active measurements to characterize the nameserver deployment of studied domains and investigate their configuration issues and potential pitfalls.

More specifically, we examine ADNS replication among the government domains to gain insights as to whether they are well maintained for high availability and reliability. We focus on evaluating the resilience of government-controlled domains' ADNS deployments both over the past 10 years and at the present. Furthermore, we investigate two fundamental configuration issues that are still common within the DNS: defective delegations and delegation inconsistency. We particularly look for risks of government domains that may cause service degradation or even service interruption.

Our major contributions are summarized as follows:

- We conduct an extensive measurement study of authoritative DNS server deployments for zones related to the e-governments among over 190 countries using both active data collection and passive DNS datasets.
- We perform an analysis based on 10 years of PDNS data and observe thousands of domains using only one nameserver each year, and a 60% increase in the countries

relying on any single third-party ADNS provider between 2011 and 2020.
- We reveal different patterns of misconfigurations, including stale records leaving over 1,000 government domains open to hijacking.

The rest of the paper is organized as follows. § II provides relevant background. We describe our approach in § III and our results in § IV. § V discusses the limitations and the implications of our findings. § VI surveys related work, and finally, § VII recapitulates our study.

II. Background

A. DNS and the Role of Nameservers

The DNS is a globally distributed database, in which the responsibility for managing portions of the database is delegated to independent entities. The system is organized into what are known as *zones*. The presence of nameserver (NS) resource records (RRs) defines zone boundaries: a domain that has an NS record is the "top node" of a zone [9]. A key concept within this framework is that of parent and child. A parent is defined as "The domain in which the Child is registered" [10]. A child is basically a subdomain of a parent domain, where both the parent and the child have NS records. A child domain is necessarily a subdomain of its parent, but not all subdomains are child zones. Rather, zones are defined by the position of NS records. Where NS records appear in the DNS, they mark the "top" of a zone. A zone file consists of the set of authoritative resource records identified within the zone.

Each zone must have one or more authoritative nameservers that are responsible for maintaining and distributing the zone's RRs, including those with information about how to find the authoritative nameservers of child zones. This role makes authoritative nameservers a vital piece of the DNS. Given the need for the high availability and reliability of these servers, a zone should have multiple authoritative nameservers [9] placed in different locations and networks [11]. Paradoxically, while authoritative nameserver replication maintains high availability and reliability, it also introduces various problems with the consistency of NS records. Further, as long as at least one authoritative nameserver of a domain can be reached and works properly, the misconfigurations or failures of other authoritative nameservers may not be detected by the domain owner in a timely manner. This type of scenario plays out relatively often in the real world and can cause problems in service availability, and create security risks.

B. Government and the DNS

ICTs have been playing an increasingly important role in supporting and shaping government operations. The government use of ICTs is commonly referred to as e-government [12]. This concept goes beyond simply augmenting existing systems with technological tools, as e-government is expected to shape the way in which citizens and governments interact, and to add value to the services governments provide [13]. While technology is not the focus of e-government, it is the foundation [14]. Thus, robust e-government will necessarily require a reliable technological foundation.

High availability, reliability, and security are essential to e-government. One of the primary ways in which e-government creates value is to enhance trust in governments [12], [13]. Trust — in technology as well as in governments themselves — also plays a key role in determining whether or not citizens will use e-government resources at all [15]. It follows that governments have a vested interest in promoting trust in the systems they have created. In this endeavor, the DNS plays an important role. Many government digital resources and services are represented by a domain name, and located using the DNS. The DNS is thus a key piece of e-government operation. In this study, we focus on ensuring DNS deployments for government domains are robust. This is an extension of the work in [16], using additional data and measurements.

III. Datasets

We constructed the datasets in three stages. First we identified the government domains to be examined. We then collected the DNS data using both a passive DNS database and active probes. Finally, we filtered the DNS data.

A. Selecting Domains

Obtaining a representative list of domain names dedicated to government use is not a straightforward task due to the diversity of how governments manage their resources. A wide variety of entities in addition to those dedicated to governance may be considered government resources (*e.g.*, universities, utilities, hospitals), but the extent to which these entities are associated with governments may vary by country, municipality, or city. Identifying and categorizing such resources to support a systematic and coherent analysis was beyond the scope of this stage of our research. Thus we focused on domains associated with national governments, as we could confidently identify these and efficiently present measurements.

To identify these domains, we used the United Nations E-Government Knowledge Base [8]. For each of the 193 UN member nations, the Knowledge Base website contains a link to the nation's designated *national portal*: a central site for e-government resources. The information about countries' e-government is partly self-reported, and UN researchers also examined the national government websites [7]. This approach to identifying government domains lends credibility to the list we obtained. That said, we did find it necessary to modify the list slightly based on additional data found via the UN site. Eleven of the links we obtained from the UN referred to domains that we could not resolve. For two of the countries involved, the registered domains in the link on the UN site differed from that in the member states questionnaire (MSQ) [7]. We also found one case where the domain in the link belongs to a third-party that is using it to serve search results and advertisements. For this case and the two involving a mismatch between the link in the page and the MSQ, we used the domain in the MSQ.

The FQDNs in the national portal links provided a starting source to build lists of domain names that allowed us to study countries' government DNS deployments in greater depth. The FQDN itself may be associated only with the national portal website, while other government-related resources exist in the same zone. Thus, for each FQDN, we extracted the suffix or the registered domain of the FQDN in the link, and used that to seed further searches. For example, given the FQDN `www.australia.gov.au`, we used the suffix `gov.au`. With this approach, we needed to ensure that a government entity controls the registered domain or the suffix. To do so, we did a manual search of the documentation for the country code top level domain's (ccTLD's) registration provider listed in IANA's Root Database [17]. We used this documentation to determine if the suffix was reserved for government use. In cases where we could find such information, we checked a registrar to see if the suffix was listed as restricted. We found only three cases (`laogov.gov.la, timor-leste.gov.tl and jis.gov.jm`) where we could not verify that the suffix was reserved for government use. In these cases, we used the registered domain rather than the suffix. Additionally, we identified only one FQDN, `www.regjeringen.no` (Norway), which has NS records but is not covered by our suffix check. In this case, we verified that the registered domain (`regjeringen.no`) is associated with the government via the MSQ and Whois information. We refer to the set of the seed domains as d_{gov}.

To grow a larger list of government-controlled domains, we used Farsight's DNSDB [18] to retrieve NS records for the d_{gov} and their subdomains. A global network of sensors and several zone files provide the input to the DNSDB [18]. The DNSDB, which has been maintained since 2010, contains over 130 billion unique record sets with data for more than 51 billion FQDNs [19]. This dataset allowed us to discover zones within the namespace defined by our seed domains. We use left hand wildcard searches in the standard DNSDB to retrieve NS records for each selected seed domain. As we intended to use these domains for active queries, we wanted to identify domains that were likely to still be in use. We noted that those reported by sensors or seen in zone files relatively recently were the most likely to fit this need. Thus, from the PDNS data, we extracted all FQDNs from records seen in the database inputs between January 1, 2020 and the time at which we collected the data in February 2021.

B. Data Collection

We used the PDNS data to generate a list of domains to examine via active lookup. After some filtering to remove what appeared to be disposable domains, we obtained a list of over 147 thousand domains to query. Given this list of domains to study, we ran a series of DNS queries, retrieving the domains' NS records. This collection was performed through a server in our university's network in April 2021. Figure 1 illustrates the measurement setup. Given a subdomain d, the client first identifies the authoritative nameservers of d's parent, which will be queried for d's NS records (①). If one of these nameservers returns a referral (②), we will proceed to the next step. In this step (③), the query client sends the same query to d's authoritative nameservers. If one of d's authoritative nameservers returns an answer (④), we will combine the authoritative nameservers returned this answer with those obtained in step (①). Finally (not shown in the figure), the client retrieves the IPv4 addresses of all authoritative nameservers identified in the previous steps and sends a query to each address for d's NS records.

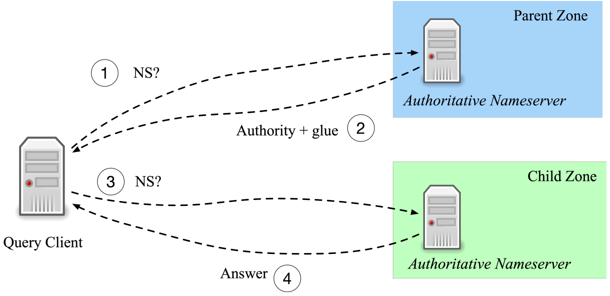

Fig. 1: Measurement setup for active data collection

For cases where the authoritative nameservers of the parent returned NS records, but the nameservers listed in those records did not reply, we ran a second round of queries in case the inability to reach the latter was due to transient conditions. The second round of queries was started shortly after the first, and the interval between subsequent queries varied per domain from a few minutes to a few days. We did not re-run queries in cases where the parent zone's authoritative nameservers did not reply, as we expected many of the domains in these cases are simply not active anymore, and querying again would create unnecessary additional traffic.

Given our efforts to make our experiments efficient, one might question why we did not simply use the DNSDB to collect data for more of our measurements. That is, why not retrieve `A` and AAAA records for the nameservers we studied and use these to infer information about replication and delegation? Such an approach has been used in other works [20], [21]. However, doing so requires certain assumptions about the coverage of the PDNS dataset that are not suitable for our study. For example, these other studies generally examine domains at the second level of the DNS hierarchy, whereas we largely study domains at lower levels. Less than 1% of the domains we examined were second-level domains. Most (85.4%) were third-level domains and 10.9% were fourth-level domains. Information for domains below the second level may not appear in zone files. Further, PDNS data could not support some of the measurements we wished to conduct, such as identifying unresponsive nameservers. Thus, we relied on the PDNS data to build our list of domains to query, and used active measures to assess the state of these domains' authoritative nameserver deployments.

Figures 2 and 3 summarize the statistics of the PDNS data. As Figure 2 shows, the number of domains with NS records seen in the data grew from 113.5 thousand in 2011 to 192.6 thousand in 2020. The slight decrease from 2019 to 2020

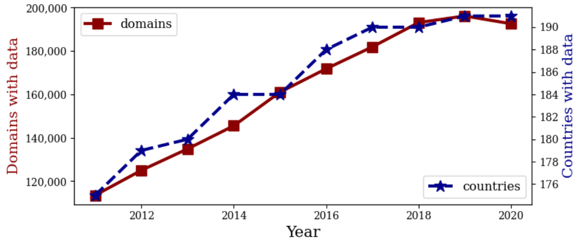

Fig. 2: Number of domains and countries with data in PDNS from 2011 to 2020

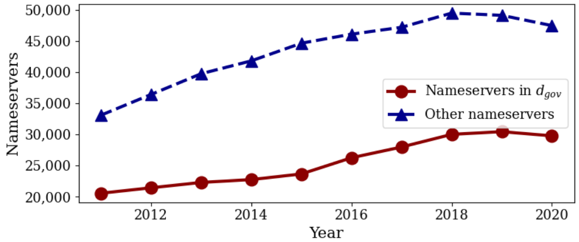

Fig. 3: Number of nameservers in PDNS from 2011 to 2020

appears to be due to a consolidation of the domains for governments at different levels in China. Figure 3 shows similar growth patterns in the number of nameservers (hostnames).

Our active data collection yielded 115 thousand domains with at least one response from a nameserver in the parent zone. For 96 thousand of domains, at least one such response was not empty. Figure 4 shows the distribution of the responsive domains by country. As this figure shows, some countries with relatively large populations and developed e-governments have relatively little data in our dataset. We consider that there are two main reasons for this. First, a country's e-government may be highly centralized and use few zones. Second, the country may use a domain other than that which is associated with the national portal for many of its e-government resources. We discuss this limitation further in V, and in the following sections we present results by country in addition to aggregated numbers.

C. Data Filtering

We filtered the PDNS data by removing records that appeared for only a short time. Such records may represent a variety of scenarios, including misconfigurations, the use of DDoS protection services, or domain expiration. We are primarily interested in characterizing stable, consistent deployment strategies, and thus we removed records in which the difference between the last-seen and first-seen timestamps was less than a minimum number of days. We set this minimum based on the maximum TTLs of a few popular resolvers [22]–[26], selecting the largest TTL, 7 days. In a scenario where an issue can be quickly detected and corrected, we expect the incorrect records could continue to show up for 7 days due to caching. Thus, we use a duration of 7 days as a threshold to differentiate between stable and transient records.

A second step in filtering the PDNS involves identifying the earliest date on which we can consider that a domain was used by a government entity. For those d_{gov} where we use a suffix, the suffix is reserved for government use. We assume that this restriction has been in effect from the time delegations within the associated zone began. Thus, while we may not have any data for the domain at the start of the period, we are unlikely to have data for a non-government entity using that domain. For other domains, we use the Web Archive [27] to find the earliest date on which a website appeared at the domain belonging to a government entity. While fraudulent websites posing as official resources are common, we consider it unlikely that a government would take over a website previously controlled by imposters and use that domain to serve legitimate content.

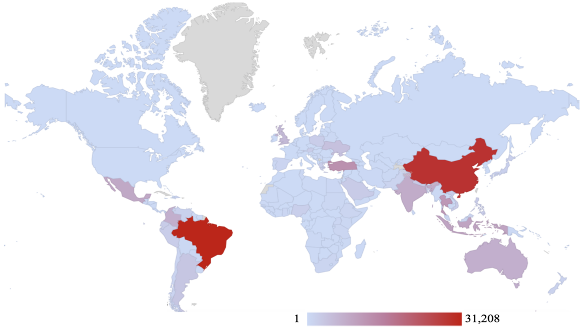

Fig. 4: Number of domains per country in PDNS data, 2020

D. Ethical Considerations

In our study, a primary concern was to ensure our measurements would not create an unreasonable load, and that operators of the domains we queried could identify and contact us. To that end, the server used to run measurements was assigned a static IP address, for which a PTR was created to indicate that the server was used for research purposes. We also limited the rate of our queries. Also, in the case of PDNS data, user privacy is the primary concern. The data we deal with has all information that might identify original clients removed [28]. Additionally, for the studied domains, we do not attempt to reconstruct zone files or map a domain's network. We have taken steps toward responsible disclosure, contacting operators of domains in which we found vulnerabilities.

IV. Characterization of Government DNS

This section presents the results of various measurements of ADNS configurations, including measurements of replication, dependency, delegation, and consistency.

A. Nameserver Replication

The number of designated authoritative nameservers is one important metric for assessing a domain's availability and reliability. Relevant RFCs require that a domain have at least 2 authoritative nameservers, and note that in many cases having more than 2 is better [9], [11]. Further, these nameservers should be in different physical locations and networks [11]. Although there is some debate as to whether replication is practical in all scenarios [29], the number of a domain's nameservers does provide a helpful perspective on the domain's

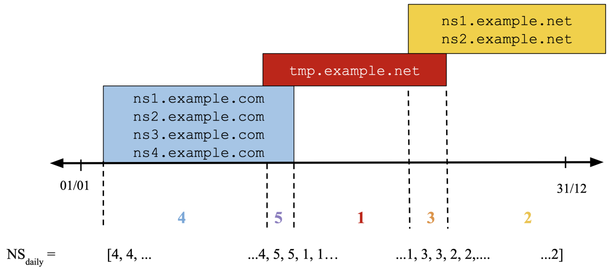

Fig. 5: Illustration of identifying the number of authoritative nameservers for a domain in a given year

ADNS deployment strategy, and is often used by researchers for examining ADNS deployments [21], [30]–[32].

In the following discussion, we refer to a domain relying on a single authoritative nameserver as d_{1NS} . We examined trends in the prevalence of such domains over the past decade using the PDNS data by identifying the deployment strategy per domain per year. That is, for a domain, d, in any given year, we first determined how many nameservers d used on each day of the year. We represented this information as a list NS_{daily}, where each element is a number that represents the number of nameservers that were in the NS records for d on each date (see Figure 5). NS_{daily} can contain up to 366 elements, or as few as 1 element (since we do not consider days on which no NS records appear to be active). We used the mode of NS_{daily} to represent the state for d for that year.

Single-nameserver Domains. The results of our measurements on the change in the prevalence of d_{1NS}s over the past decade inform us of a mixed story. Between 2011 and 2020, the total number of d_{1NS}s increased, but at a lower rate than the overall number of domains in the dataset. The former increases by a factor of 1.2 (from 4.8 thousand to 5.9 thousand), and the latter by a factor of 1.7 (from 113.5 thousand to 192.6 thousand). For most countries, the number of d_{1NS}s decreases (34) or remains the same (98). In most cases (92), those countries experiencing no change had no d_{1NS}. The increase in the percent of domains using at least 2 nameservers seems to indicate a trend towards increased replication. However, as the total number of d_{1NS}s also appears to have increased, this is clearly a persistent pattern. To understand this pattern, we examined the d_{1NS}s further.

For each year, in the range [2012, 2020], we computed the percentage of the d_{1NS}s that were new, and the percentage of that were observed to be using a single nameserver in 2011. We also found what percent of d_{1NS}s from 2011 were no longer active. As shown in Figure 6, the overlap decreases steadily, and by 2020, only 21% of the d_{1NS} from 2011 were still active. Measuring the overlap between sequential years yielded similar results. The percentage of all d_{1NS}s that were new in a given year ranged 14%-23%, and the percentage that were no longer seen was 16%-26%. This consistent change suggests that the pattern of d_{1NS}s cannot be attributed to a single group of domains that exist across the years, but to persistent patterns in authoritative NS deployments. We

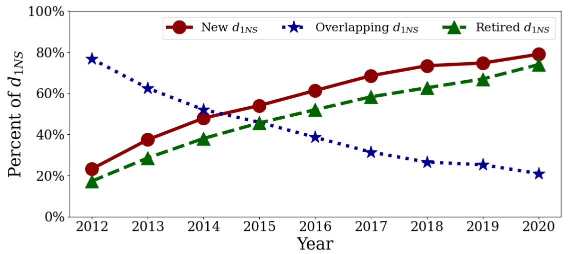

Fig. 6: Changes in set of of all d_{1NS} per year

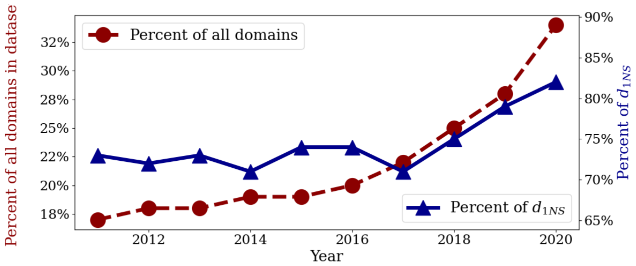

Fig. 7: The percentage of D_{1NS} and of all domains using a private ADNS deployment per year

also examined what percent of the d_{1NS}s each year were using a private ADNS deployment strategy. We considered a domain to be using a private deployment if the nameserver hostname was in the same d_{gov} as that domain. As the governments we examined may operate nameservers in other domains, our measurement represents a lower bound on the private deployment. As shown in Figure 7, the percentage of d_{1NS}s using a private ADNS deployment each year was over 71%r, while the percentage of domains overall using a such a deployment was less than 34%. Investigating some of these cases suggested some d_{1NS}s belonged to relatively small entities. For such domains, the resilience gained by having multiple nameservers may not merit the effort required to operate them, or the security risks involved in using a third-party provider. We investigate the question of third-party providers further in the next section.

We considered what insights we could obtain into current deployments using our active data collection. Figure 9 provides the distribution of authoritative nameservers among the domains we studied. Of all the domains considered, 98.4% used at least two authoritative nameservers, and for over half (109) of the countries considered, no domain in its d_{gov} used less than two nameservers. In contrast, for 15 countries, at least 10% of responsive the domains used a single authoritative nameserver. Four of these (Bolivia, Bulgaria, Burkina Faso, and the UAE) had fewer than 10 responsive domains of which only a few (three or fewer) were d_{1NS}. Of the other 11, in three cases (Indonesia, Kyrgyzstan, and Mexico), over half of the d_{1NS} had no response from their authoritative nameservers, suggesting these domains are no longer in use, but have not been removed from parent zones.

Using active DNS lookups, we examined how common it is for d_{1NS} to represent stale records. We consider that if we could not obtain a response from a domain's authoritative nameservers, the domain is no longer used and the NS records

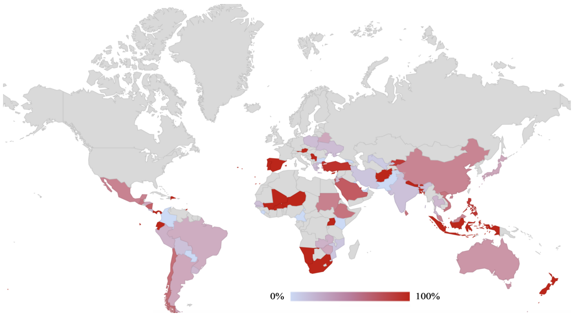

Fig. 8: Percentage of d_{1NS}s with no authoritative response from authoritative nameservers.

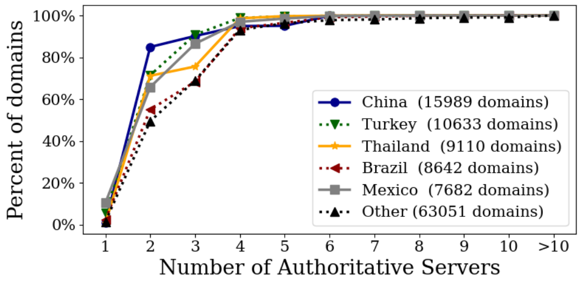

Fig. 9: CDF of the number of ADNS listed in NS records

in the parent zone are stale. Note that there may be cases in which a domain is no longer used by its owners but a nameserver still provides a response. Of the d_{1NS} we found in our active measurements, 60.1% had no response from an authoritative nameserver. This percentage is much higher for d_{1NS} under particular d_{gov}, as shown in Figure 8.

Diversity of Nameservers. Not just the number of physical hosts, but also the placement of authoritative nameservers plays a key role in maintaining a robust ADNS deployment. Using the active measurements, we evaluated the domains with multiple nameservers. For each domain, we identified the set of IPv4 addresses (IP_{ns}) to which its nameservers resolve. Using Maxmind's GeoIP2 ASN (Autonomous System Number) database [33], we determined the number of /24 prefixes (24_{ns}) and ASNs (ASN_{ns}) to which the IPs in IP_{ns} belong. Table I shows these statistics. Curiously, we observed some cases of a single IP address being used, even when the domain has multiple nameservers listed. More than half of these cases were linked to domains in a single d_{gov}. Many of these domains are sharing nameserver pairs that resolve to the same IP address. Regarding IP diversity, in most cases the nameservers for a given domain cover multiple /24 prefixes, although less than a third cover multiple autonomous systems. While even nameservers in multiple prefixes may still share a single-point-failure (*e.g.*, a border router), our results suggest that overall, most zones have a diverse placement of authoritative nameservers.

Whether we use the /24 prefix or autonomous system to estimate replication, the percentage of domains with sufficiently distributed authoritative nameservers is relatively low compared to the results in [21]. In that work, the author found over 85% of domains studied used authoritative nameservers with IP addresses in multiple /24 prefixes, while only 72% of domains actually had servers in different networks. The discrepancy between the results in this work and ours may be attributed to the fact that the other work focused on popular domains (the Alexa Top 1 Million), and domains at the second level of the DNS hierarchy. Intuitively, popular domains and those higher in the hierarchy require more robust ADNS deployments than do less popular domains or those lower in the hierarchy. We see in our own results that the percent of domains with authoritative nameservers resolving to IPs in multiple /24 prefixes is higher for domains at the second level of the DNS hierarchy (87.1%) than for those at the third through fifth levels (less than 80%). It is difficult to separate patterns for domains at different levels from those of different governments, though. Delegation strategies cause some countries to dominate the set of domains in our dataset at different levels. For example, 16% of the domains in our dataset at the third level of our DNS hierarchy were in `gov.cn`, and 53% of those at the fourth level were in `gov.br`. Domains in different countries tended to have different deployment styles or use different providers. For the responsive subdomains of `gov.cn`, over half were using authoritative nameservers under `hichina.com` (38%), `xincache.com` (19%) and `dns-diy.com` (10.8%). In contrast, for domains in `gov.br` the maximum percent of domains using any given provider was much smaller: only 6% using authoritative nameservers belonging to Hostgator. Dissecting patterns in ADNS deployment strategies further is beyond the scope of this work. Overall, we observe high levels of replication, although the percent of authoritative nameservers per domain in different networks appears to be relatively low compared to that of popular domains.

TABLE I: The number of IPv4 addresses, /24 prefixes and ASNs associated with domains with multiple nameservers

	Domains	$\|IP_{ns}\|$ >1	$\|24_{ns}\|$ >1	$\|ASN_{ns}\|$ >1
Total	94,848	89.8%	71.5%	32.9%
China	13,623	97.3%	95.7%	52.4%
Thailand	8,941	36.1%	31.7%	13.6%
Brazil	7,271	95.7%	54.4%	13.7%
Mexico	5,256	90.0%	67.4%	25.7%
UK	4,788	99.7%	96.1%	25.5%
Turkey	4,528	91.1%	72.6%	42.1%
India	4,426	93.4%	84.1%	10.6%
Australia	3,707	99.2%	91.7%	9.0%
Ukraine	3,421	99.0%	62.3%	45.1%
Argentina	2,795	97.6%	71.8%	30.5%

B. Third-Party DNS Providers

As more domains have their authoritative nameservers managed by third-party DNS service providers, the degree to which these domains depend on such providers has become an increasingly important question to answer. Questions regarding DNS providers are, in a way, an extension of those regarding nameserver replication. A single domain may have multiple nameservers in diverse locations, but experience has shown

TABLE II: Government usage of major DNS providers (ordered alphabetically)

Provider	Domains	d_{1P}	Sub-Regions	Domains	d_{1P}	Sub-Regions
	2011			**2020**		
Amazon	5 (0.0%)	1 (0.0%)	3 (9.4%)	5193 (2.7%)	4712 (3.4%)	27 (84.4%)
Azure	0	0	0	1574 (0.8%)	1155 (0.8%)	24 (75.0%)
Cloudflare	12 (0.0%)	5 (0.0%)	6 (18.8%)	4136 (2.1%)	3104 (2.3%)	31 (96.9%)
DNSPod	373 (0.3%)	181 (0.3%)	1 (3.1%)	700 (0.4%)	575 (0.4%)	1 (3.1%)
DNSMadeEasy	89 (0.1%)	50 (0.1%)	13 (40.6%)	254 (0.1%)	220 (0.2%)	16 (50.0%)
Dyn	7 (0.0%)	1 (0.0%)	3 (9.4%)	170 (0.1%)	131 (0.1%)	13 (40.6%)
GoDaddy	283 (0.3%)	190 (0.3%)	19 (59.4%)	1582 (0.8%)	1262 (0.9%)	20 (62.5%)
UltraDNS	15 (0.0%)	5 (0.0%)	4 (12.5%)	66 (0.0%)	57 (0.0%)	6 (18.8%)

* In Tables II and III, groups refer to all countries in a sub-region, with the exception of the top 10 countries with the most records in the PDNS data. Sub-regions are defined by the UN [8].

TABLE III: Top DNS providers ranked by the number of countries with subdomains using the provider

Provider	Domains	Sub-Regions	Countries	Provider	Domains	Sub-Regions	Countries
	2011				**2020**		
websitewelcome.com	424 (0.4%)	23 (71.9%)	52.0	cloudflare.com	4,136 (2.1%)	31 (96.9%)	85.0
domaincontrol.com	283 (0.3%)	19 (59.4%)	47.0	AWS DNS*	5,193 (2.7%)	27 (84.4%)	67.0
zoneedit.com	182 (0.2%)	21 (65.6%)	32.0	domaincontrol.com	1,582 (0.8%)	20 (62.5%)	63.0
dreamhost.com	243 (0.2%)	18 (56.2%)	29.0	bluehost.com	432 (0.2%)	21 (65.6%)	58.0
bluehost.com	134 (0.1%)	14 (43.8%)	29.0	Hostgator*	1,536 (0.8%)	21 (65.6%)	55.0
hostgator	183 (0.2%)	18 (56.2%)	29.0	websitewelcome.com	745 (0.4%)	18 (56.2%)	50.0
ixwebhosting.com	98 (0.1%)	16 (50.0%)	28.0	digitalocean.com	429 (0.2%)	19 (59.4%)	45.0
hostmonster.com	103 (0.1%)	16 (50.0%)	27.0	microsoftonline.com	135 (0.1%)	20 (62.5%)	41.0
everydns.net	259 (0.2%)	17 (53.1%)	26.0	Azure DNS*	1,574 (0.8%)	24 (75.0%)	37.0
pipedns.com	48 (0.0%)	14 (43.8%)	24.0	wixdns.net	324 (0.2%)	20 (62.5%)	36.0
stabletransit.com	57 (0.1%)	13 (40.6%)	22.0	cloudns.net	225 (0.1%)	19 (59.4%)	36.0

* Nameserver domains for Amazon, Hostgator, and AzureDNS are grouped together as described in Section IV-B.

that if these are hosted by a single third-party provider, that provider could become a single point of failure for that domain. For example, in 2016, a DDoS attack targeting Dyn seriously disrupted services for Dyn's customer domains [34]. Although the risk of such providers sustaining an outage or attack that is sufficient to make a domain unavailable might be low, the past few years have witnessed other ways in which such dependency can become a real issue. It is thus of interest to assess how much influence major third-party DNS providers have upon governments. In general, a domain's nameservers are key to maintaining availability and reliability, and it is helpful to understand the general characteristics of these nameservers.

To evaluate trends in the use of DNS providers, we relied on the PDNS data, as it gave insights into both the current state and the patterns over the past few years. To examine the popularity of a provider, we needed to check what nameserver hostnames are associated with that provider. This is particularly important for major providers, such as Amazon, which use hundreds of different nameservers. For Amazon, which follows a unique naming pattern, we can identify nameservers by using a regex match. For other popular providers, we can use a combination of string-matching on nameserver domain names themselves and on the *MNAME* and *RNAME* fields in their SOA records.

Popular Third-Party Providers. We first attempted to answer the question of to what extent third-party providers are commonly used by popular domains [35], [36] (*e.g.*, the Alexa Top 1 Million) are also used by governments. We considered what percent of all domains in our dataset used these providers in 2011 and 2020, and what percent of the countries had at least one domain using these providers — that is, how widespread was their usage. To gain insight into regional trends, we grouped countries using the sub-region assigned by the UN, and measured how many of these sub-regions contained countries whose e-government domains relied on top providers. Note that for the 10 countries with the most records in the PDNS data, we considered these as having unique behavior, and treated each of these as its own sub-region. We also examined what percent of domains relying on a single provider (d_{1P}) used one of these popular providers.

As Table II shows, the use of these major providers is global, and their reach is growing. We notice that in particular the percent of domains using Amazon and Cloudflare increased slowly but steadily over the past several years from virtually none to more than 2%. While this may seem to be a small change, the number of domains using these providers increased by multiple orders of magnitude, and many of these domains are only using nameservers belonging to these providers. This pattern is consistent with the observations in other works reporting centralization within the DNS [35], [36]. Even though many governments use their own DNS infrastructures or DNS services of local providers, the concentration of domains using these top providers has been increasing. If this trend continues, the concerns of increasing centralization in DNS nameserver deployment will apply to government domains.

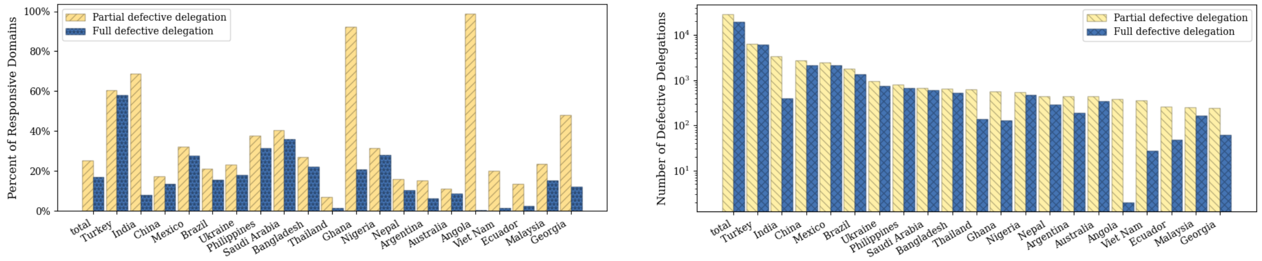

(a) Percentage of domains per country with a defective delegation involving a nameserver in P

(b) Percentage of d^s_{gov} per country with a defective delegation involving a nameserver in P

Fig. 10: Percentage and number of domains per country with an unresponsive nameserver for 20 countries with the highest number of defective delegations.

The previous measurement was guided by a list of DNS providers used commonly by popular domains. It may be that other providers have a greater impact among government domains. To check if this was the case, we identified the top 10 providers, ranked by the number of countries served, in 2011 and 2020. Table III shows the prevalence of several providers other than those indicated as common among popular domains, though we observed again the rise of Cloudflare and Amazon. We also noticed that the top 10 providers in 2020 account for a larger portion of the domains in the dataset than in 2011. Meanwhile, the number of countries with domains using any single provider grew 60% from 52 in 2011 to 85 in 2020. This again highlights increased centralization, although at this point, the DNS ecosystem of government domains is still highly heterogeneous.

C. Defective Delegations

A defective delegation (usually called a *lame delegation*) occurs when a nameserver included in NS records for a zone does not answer queries for that zone. Defective delegations have several causes, including configuration errors, changes in nameservers without an update to the parent zone, and changes in the services of a third-party provider. In some cases, defective delegations pose a serious security risk, leaving a domain vulnerable to hijacking or monitoring. Even in cases where the risk of monitoring or hijacking is low, defective delegations can lead to performance degradation, due to increased latency and extra traffic [1], [20].

Defective Delegation Prevalence. Using the terminology from [20], we refer to cases in which none of the authoritative nameservers listed for a domain are able to provide answers for that domain as a *fully* defective delegation. Cases where at least one authoritative nameserver does not respond, we refer to as *partially* defective defective delegations. Note that fully defective delegations are a subset of the partially defective delegations. We examined both partially and fully defective delegations by checking the data from our active lookups. Figures 10a and 10b summarize the major patterns we observed in defective delegations. Surprisingly, 29.5% of the domains had a defective delegation. Slightly more than a quarter of the domains (25.4%) had a partial defective delegation involving the information just in the parent zone. This pattern is largely driven by a few d_{gov} with a large number of subdomains and a relatively high rate of defective delegations. A few countries have many more partially defective delegations than fully defective delegations. In these cases, most of the domains involved were sharing an authoritative nameserver that either could not be resolved or was no longer serving those domains.

Hijacking Risks. We also explored what percentage of these defective delegations would present a security risk for domain hijacking. Most of them presented little to no risk of a hijacking, as they involved nameservers belonging to governments themselves. We checked how many of the nameserver domains not in a government domain were available, and found 805 that could be registered (using GoDaddy). The cost per domain ranged from 0.01 to 20,000 USD, with a median of 11.99 USD (see Figure 12). These were used by 1,121 domains in 49 countries. Figure 11 shows the number of affected domains per country, and the number of available nameserver domains for those countries with the most affected domains. Only 2 available nameserver domains were used by domains associated with governments of more than one country. Similarly, for almost one third of the countries whose subdomains had defective delegations, these delegations pointed to nameservers in a single domain. Several of the nameservers involved in the defective delegations had patterns suggestive of typos. For example, the parent zone for one domain listed `pns12cloudns.net`, `pns13cloudns.net`, and `pns14cloudns.net`, along with `pns11.cloudns.net`, which appears to be the only functioning nameserver of the four. For more than half (625) of the domains in this scenario, we did not receive any authoritative response from the authoritative nameservers at all, suggesting the nameservers of these domains are no longer active. For the d_{gov} with the most affected subdomains, including Turkey, Brazil, and Mexico, the majority of the subdomains were in this group. These appear to be cases of stale records, in some cases, dozens or even hundreds in the same d_{gov}.

The existence of dangling NS records [37], though not particularly surprising given the demonstrated prevalence of this type of scenario in other areas, is nonetheless a serious concern. The ability to control the resolution of domains with the same suffix domain as legitimate government domains provides an avenue for serious attacks.

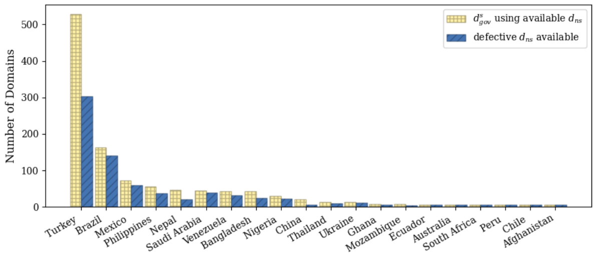

Fig. 11: Available d_{ns} in defective delegations, by country

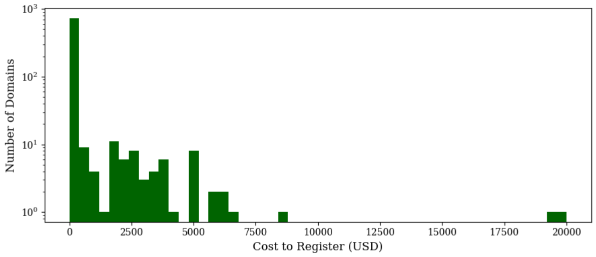

Fig. 12: Distribution of the cost to register available d_{ns} in defective delegations

D. Parent-Child Inconsistency

The specifications for the DNS require that parent and child zones should remain consistent [9]. That is, all of a zone's authoritative nameservers should contain the same NS RRs for any given child zone as those contained in the child zone's authoritative nameservers. Inconsistency between zones could result in unexpected load distribution patterns and increased latency [20], [38]. In some cases, the inconsistency may also indicate stale records and, as with defective delegations, the potential for domain hijacking and privacy leakage.

To characterize the inconsistency between parent and child zones, we followed the framework used in [39]. According to this approach, for each domain we queried, we first checked if the nameservers listed in the parent zone (P) and those listed by the domain's authoritative nameservers (C) were identical. If they were not, we checked if the two sets have at least one authoritative nameserver in common. When no intersection existed, we explored to what extent there was an intersection between the IPv4 addresses to which the authoritative nameservers P and C resolve (denoted as $IP(P)$ and $IP(C)$ respectively). If P and C overlapped by at least one authoritative nameserver, we checked to see if P included all authoritative nameservers in C or vice versa, or if neither case held.

Inconsistency Prevalence. Figure 13 summarizes our findings of zone inconsistency. For most (76.8%) of the domains that were responsive, no inconsistency between the authoritative nameservers in P and C appeared. While this percentage was fairly high, it was substantially lower than that in [39], in which the percentage of cases where $P = C$ was more than 90% of responsive domains in all the zones studied. We found that the level of consistency is much higher (93.5%) for domains at the second level of the DNS hierarchy than domains with three or more levels (77% or less). In [39], only domains at the second level domains are considered. We would expect more inconsistencies between zones further down in the hierarchy, as these will generally be operated by increasingly smaller entities that may have fewer resources to devote to maintaining their DNS records, and may not experience as many problems due to zone inconsistency as a larger entity would. This speculation can be confirmed by the observation in [38], which reports that inconsistencies are more common below the second level of the `.edu` domain than in the second level itself.

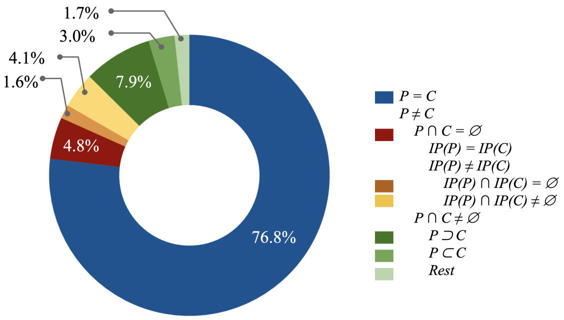

Fig. 13: Summary of parent/child zone consistency

As in the case with d_{1NS}s, the prevalence of disagreement between zones varied widely by country, as shown in Figure 14. The countries with the largest percent of domains having a disagreement tend to have few responsive domains, but there are also some countries with a large number of responsive domains where this behavior is relatively common.

Inconsistency between a parent and its child zone may represent misconfigurations or stale NS records. For example, we observed that in several cases where $P \neq C$, at least one of the authoritative nameservers involved is not a fully qualified domain name, *i.e.*, a single label such as *ns* or *dns-server*. This type of error arises from typos in zone files where a trailing '.' is added to what should be a relative domain name. In this scenario, the authoritative nameserver cannot append the origin to the name. As an illustration, consider an NS record for a domain using `ns.example.com` as an authoritative nameserver. If the record has the entry *ns.* rather than *ns*, the authoritative nameserver returning the NS record will simply send *ns* rather than `ns.example.com`.

Hijacking Risks. We found that 40.9% of domains for which $P \neq C$ also have at least one partially defective delegation. The previous section explores these cases. Furthermore, a dangling NS record may also exist even when there is no defective delegation. For example, if the authoritative nameserver indicated in the parent zone now belongs to a parking service, the server may respond to all DNS queries with answers directing users to their own servers. To explore this scenario, we checked the cases of inconsistency where the authoritative nameservers involved were not defective. As with the defective delegations, we identified the domain names of the authoritative nameservers that were not included in both P and C, and checked to see if any is available for registration. We found 13 d_{ns} that were available for registration. In

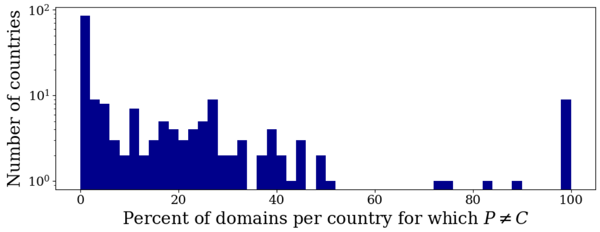

Fig. 14: Distribution of the rate of disagreement per d_{gov}

our dataset, these d_{ns} served 26 domains in seven different countries. Twelve of the domains, representing district governments in one country, had evidently all used the same DNS provider, and that provider's domain was expired. Other cases were related to various entities, including a ministry of finance, a group responsible for taxes, and local governments. The minimum cost to register any of the d_{ns} with GoDaddy was 300 USD. As in the case with available domains found through defective delegations, the cost to leverage one of these dangling records is not high, and may easily merit the reward for some adversaries.

V. Discussion

Several of the limitations and assumptions of this work have been mentioned previously in the context of the relevant measurements. We reiterate some here for clarity. We also discuss what we consider promising approaches to addressing the issues we have highlighted.

A. Limitations

Some governments may use domains besides the one associated with the national portal for many of their resources. By examining each country individually, we might have been able to identify these domains. However, language and cultural variations, as well as the presence of sites spoofing government resources would complicate such a search. Including domains that we could confirm with confidence would introduce additional bias. Thus, at this stage of the work, we kept a relatively narrow, objectively-defined scope. We will further expand the dataset in our future work.

For the active measurements, our data were collected from a single vantage point in the United States. It is possible that we might obtain different results if using additional vantage points. Such a situation could arise if the authoritative nameservers we studied only answer queries from a specific range of IP addresses, or return different responses based on the IP addresses from which queries are sent. However, since e-government websites are unlikely to distribute content from multiple sites and use geolocation-based content, we do not anticipate the results would vary greatly across multiple vantage points. Thus, conducting measurements from additional vantage points will be an interesting direction for future work, but with a lower priority.

As noted in Section III-C, we repeated certain measurements to account for transient failures. However, we did not repeat all that we might have. Specifically, we re-ran queries for a domain if at least one nameserver in its parent zone returned NS records for the domain but no authoritative nameserver sent a response. However, if a nameserver was unresponsive while other authoritative nameservers in the same zone did send answers, we did not re-send queries to the unresponsive nameserver. This may have led us to moderately overestimate the number of defective delegations. However, since many authoritative nameservers are used by multiple domains, most were checked at least twice. We found that only 5.7% of the authoritative nameservers we examined cannot be resolved *and* were only checked once, affecting 2,424 domains (less than 2% of the domains queried).

We examined the risks of domain hijacking, but did not attempt to determine if any such attacks have taken place. We could have explored it with active measurements or via the PDNS data. However, verifying a domain hijacking attack, particularly when using historical data, presents several challenges, as domain owners may periodically change the infrastructure they use. We will further investigate this problem in our future work.

B. Potential Remedies

Addressing the issues we have explored is not a simple task. Misconfigurations have plagued the DNS ecosystem for decades, despite the availability of several tools to detect or correct them. Vulnerabilities such as those we have highlighted are hardly even surprising but nonetheless demand attention.

Regarding the number and diversity of authoritative nameservers, zone operators may be challenged to find a balance between redundancy and dependency. While there may be inexpensive options for achieving replication, these can introduce new risks. For example, the use of third-party DNS providers increases the attack surface of a domain [40], and attackers can leverage vulnerabilities among such providers to hijack domains [2]. Using such providers might also tend to increase centralization within the DNS. Overall, updated guidance or requirements on developing ADNS deployments would be helpful to address this situation.

The problems of defective delegation and inconsistencies between zones are different matters, as the approaches to correct these are more well-defined. Since the inception of the DNS, various groups have developed tools for DNS debugging [41]–[44]. Also, some popular authoritative nameserver software has the capability of detecting defective delegation or inconsistency between zones [22], [45], [46]. Additionally, methods to streamline synchronization between various parties in the DNS also exist. For example, the Extensible Provisioning Protocol allows registrars to interact with registries in an automated fashion [47]. Also, the CSYNC record type provides a way for authoritative nameservers in parent zones to automatically retrieve updates from their child zones [10]. However, these tools are not without their own complications. Specifications for EPP and CSYNC processing include provisions that require out-of-band communications for certain updates. For example, depending on how the *immediate* bit in a CSYNC RR is set, the party responsible for updating the parent zone may or must require the child zone operator to confirm changes out-of-band [10]. These provisions defend against DNS hijacking. Indeed, to combat such attacks, some researchers recommend implementing defensive measures such as registry locks that would explicitly require human interactions to change a domain's authoritative nameservers [48], [49].

VI. Related Work

The previous works pertinent to our research include those studying ADNS robustness, centralization, and misconfigurations, as well as those specifically focused on government or regional DNS resources.

A. DNS Deployment Strategies

Several studies have evaluated the robustness of DNS authoritative server deployments, often using replication as a key metric. In an early work (1992), Danzig *et al.* [32] examined this aspect of deployments for second- and third-level domains in the *.edu* namespace. Callahan *et al.* [31] conducted measurements of authoritative server replication when evaluating DNS traffic captured in a residential network for 14 months between 2011 and 2012. Hao *et al.* [30] elaborated on measurements of redundancy by characterizing DNS deployment strategies and evaluating redundancy for domains using different strategies. Their work focused on ADNS deployments for the Alexa top 1 million, and data was collected in 2014 and 2015. In a study covering a 9-year period (2009-2018), Allman [21] evaluated trends in the number and topological diversity of authoritative servers for domains' in the *.net*, *.org* and *.com* domains.

From a different perspective, researchers have also examined various common misconfigurations in ADNS deployments. Pappas *et al.* [50] highlighted defective delegations and cyclic dependencies. Their measurements were conducted using passive DNS from a university campus and active measurements for domains randomly sampled from reverse zone files. They demonstrated the negative impacts of misconfigurations on performance within the DNS. Kalafut *et al.* [51] measured the prevalence and causes of cases where an authoritative server appearing in NS records with corresponding glue A does not actually exist, *i.e.*, it cannot be resolved. The authors studied the *.asia*, *.com*, *.info*, *.mobi*, *.net*, and *.org* zones. Phokeer *et al.* [52] studied the prevalence of defective delegations for reverse domains in AFRINIC. Sommese *et al.* [39] measured inconsistency between parent and child zones, and showed how these inconsistent records could negatively impact load distribution and increase latency for DNS resolutions. Their study examined the *.com*, *.org* and *.net* and root zones using data captured in 2019. Akiwate *et al.* [20] examined similar problems using 10 years' worth of zone files for several TLDs. Issues such as zone inconsistency or defective delegation often indicate the presence of dangling DNS records, which form the focus of other works [37], [53]. They may also indicate typos, and [54] examines the threat posed by such cases.

B. Region-Specific DNS Research

Also of interest are previous studies with a focus on regional DNS. In a 2012 study, Kagwe and Muthoni examined 2000 *.ke* domains [55]. In [56], the authors reported some of the challenges to DNS resilience for small island nations. In 2007, Shi [57] examined to what extent domains belonging to government entities in China conformed to the China Government Domain Name Standard. The study covered domains for 316 provincial governments and prefectures listed in 2005 and 2006. In 2008, Islam [58] examined domain names belonging to sites operated by the Bangladeshi government. These works highlight the unique challenges and patterns for various countries or geographic regions. However, none of them covers the same set of research questions we attempt to answer, making ours the first to study government ADNS deployment across a wide range of countries.

Previous research has also investigated the use of protocols designed for the security of the DNS, focusing on regional or country-specific patterns, sometimes on governments. In [59], the authors examined how registrar policies affect DNSSEC deployment, focusing on a few generic TLDs as well as the *.se* and *.nl* ccTLDs. Visoottiviseth and Poonsiri also examined DNSSEC, focusing on deployment in Thailand [60]. Similarly, in 2020 researchers examined the deployment of CAA records for governments around the world [61].

VII. Conclusions

In this work, we have shown that DNS misconfigurations persist, even among critically important domains, such as those maintained by governments. In particular, we found that while the vast majority of domains had replicated nameservers, less than three fourths had their nameservers located in different networks, and less than a third had them in different autonomous systems. We noticed an increase of multiple orders of magnitude in the number of domains relying on providers such as Cloudflare and Amazon, highlighting the trend towards dependence on a few providers, even among government domains. Finally, we uncovered defective delegations for more than 29% of the domains studied and disagreement between zones for more than 76%. Hundreds of these cases are associated with dangling records that could be exploited for domain hijacking. The prevalence of these errors among government domains is of considerable concern since these domains typically serve as authentic sources for citizens. Any failures or compromise may undermine the trustworthiness of the digital resources provided by governments. We hope that our study will promote the awareness of DNS robustness and provide insight for addressing DNS misconfigurations in the future.

Acknowledgments

We would like to thank our shepherd, Marc Dacier, and the anonymous reviewers for their detailed and insightful comments. This work is partially supported by an NSF Graduate Research Fellowship Grant No. 1247394 and an Internet Freedom Fund from Open Technology Fund (OTF).

References

[1] D. Barr, "Common DNS Operational and Configuration Errors," RFC 1912, 1996.

[2] D. Adamitis and P. Rascagneres, "Sea Turtle keeps on swimming," 2019. [Online]. Available: https://blog.talosintelligence.com/2019/07/sea-turtle-keeps-on-swimming.html

[3] "KrebsOnSecurity Hit With Record DDoS." [Online]. Available: https://krebsonsecurity.com/2016/09/krebsonsecurity-hit-with-record-ddos/

[4] N. Biasini and J. Esler, "Threat Spotlight: Angler Lurking in the Domain Shadows," 2015.

[5] "Bomb Threat, Sextortion Spammers Abused Weakness at GoDaddy.com." [Online]. Available: https://krebsonsecurity.com/2019/01/bomb-threat-sextortion-spammers-abused-weakness-at-godaddy-com/

[6] R. Houser, S. Hao, Z. Li, D. Liu, C. Cotton, and H. Wang, "A Comprehensive Measurement-based Investigation of DNS Hijacking," in *SRDS*, 2021.

[7] United-Nations, "UN E-Government Survey 2020." [Online]. Available: https://publicadministration.un.org/egovkb/en-us/Reports/UN-E-Government-Survey-2020

[8] ——, "E-government knowledgebase country data." [Online]. Available: https://publicadministration.un.org/egovkb/Data-Center

[9] P. Mockapetris, "Domain Names - Concepts and Facilities," RFC 1034, 1987.

[10] W. Hardaker, "Child-to-Parent Synchronization in DNS," RFC 7477, 2015.

[11] R. Elz, R. Bush, S. Bradner, and M. Patton, "Selection and Operation of Secondary DNS Servers," RFC 2182, 1997.

[12] OECD, "The Case for e-government," 2003. [Online]. Available: https://www.oecd-ilibrary.org/content/paper/budget-v3-art5-en

[13] J. Twizeyimana and A. Andersson, "The public value of E-Government – A literature review," *Government Information Quarterly, 36(2)*, 2019.

[14] L. T. H. Jeremy Rose, John Stouby Persson, "How e-Government Managers Prioritise Rival Value Positions: The Efficiency Imperative," *Information Polity*, vol. 20, no. 1, 2015.

[15] S. Hofmann, M. Räckers, and J. Becker, "Identifying factors of e-government acceptance–a literature review," in *International Conference on Information Systems*, 2012.

[16] R. Houser, "Investigations of the Security and Privacy of the Domain Name System," Ph.D. dissertation, University of Delaware, 2021.

[17] IANA, "Root zone database," https://www.iana.org/domains/root/db.

[18] Farsight, "Passive dns historical internet database: Farsight dnsdb." [Online]. Available: https://www.farsightsecurity.com/solutions/dnsdb/

[19] J. S. Sauver and P. Foremski, "A Decade of Passive DNS: a Snapshot of Top-Level Domain Traffic," 2021. [Online]. Available: https://info.farsightsecurity.com/a-decade-of-passive-dns

[20] G. Akiwate, M. Jonker, R. Sommese, I. Foster, G. M. Voelker, S. Savage, and K. Claffy, "Unresolved Issues: Prevalence, Persistence, and Perils of Lame Delegations," in *ACM IMC*, 2020.

[21] M. Allman, "Comments on DNS Robustness," in *ACM IMC*, 2018.

[22] ISC, "Bind 9 administrator reference manual." [Online]. Available: https://downloads.isc.org/isc/bind9/9.17.3/doc/arm/Bv9ARM.pdf

[23] N. Labs, "Unbound." [Online]. Available: https://nlnetlabs.nl/documentation/unbound/unbound.conf/

[24] S. Trenholme, "Maradns." [Online]. Available: https://maradns.samiam.org/deadwood/doc/Deadwood.html

[25] Microsoft, "Set-dnsservercache." [Online]. Available: https://docs.microsoft.com/en-us/powershell/module/dnsserver/set-dnsservercache

[26] Google, "Google Public DNS FAQ." [Online]. Available: https://developers.google.com/speed/public-dns/faq

[27] Internet-Archive, "Waybackmachine," http://web.archive.org.

[28] B. April, "Farsight DNSDB: Mapping Privacy in the DNS," https://www.farsightsecurity.com/blog/long-view/privacy-20210128/.

[29] D. J. Bernstein, "Costs and benefits of third-party DNS service." [Online]. Available: https://cr.yp.to/djbdns/third-party.html

[30] S. Hao, H. Wang, A. Stavrou, and E. Smirni, "On the DNS Deployment of Modern Web Services," in *IEEE ICNP*, 2015.

[31] T. Callahan, M. Allman, and M. Rabinovich, "On Modern DNS Behavior and Properties," *ACM SIGCOMM Comput. Commun. Rev., 43(3)*, 2013.

[32] P. B. Danzig, K. Obraczka, and A. Kumar, "An Analysis of Wide-Area Name Server Traffic: A Study of the Internet Domain Name System," in *ACM SIGCOMM*, 1992.

[34] "DDoS on Dyn Impacts Twitter, Spotify, Reddit." [Online]. Available: https://krebsonsecurity.com/2016/10/ddos-on-dyn-impacts-twitter-spotify-reddit/

[33] "Maxmind GeoIP2 Databases," https://www.maxmind.com/en/geoip2-databases.

[35] A. Kashaf, V. Sekar, and Y. Agarwal, "Analyzing Third Party Service Dependencies in Modern Web Services: Have We Learned from the Mirai-Dyn Incident?" in *ACM IMC*, 2020.

[36] M. Dell'Amico, L. Bilge, A. Kayyoor, P. Efstathopoulos, and P.-A. Vervier, "Lean On Me: Mining Internet Service Dependencies From Large-Scale DNS Data," in *ACSAC*, 2017.

[37] D. Liu, S. Hao, and H. Wang, "All Your DNS Records Point to Us: Understanding the Security Threats of Dangling DNS Records," in *ACM CCS*, 2016.

[38] J. Kristoff, "DNS inconsistency," Aug. 2018. [Online]. Available: https://blog.apnic.net/2018/08/29/dns-inconsistency/

[39] R. Sommese, G. C. M. Moura, M. Jonker, R. van Rijswijk-Deij, A. Dainotti, K. C. Claffy, and A. Sperotto, "When Parents and Children Disagree: Diving into DNS Delegation Inconsistency," in *PAM*, 2020.

[40] V. Ramasubramanian and E. G. Sirer, "Perils of Transitive Trust in the Domain Name System," in *ACM IMC*, 2005.

[41] A. Romao, "Tools for DNS debugging," RFC 1713, 1994.

[42] V. Pappas, P. Fältström, D. Massey, and L. Zhang, "Distributed DNS Troubleshooting," in *ACM SIGCOMM Workshop on Network Troubleshooting (NetT)*, 2004.

[43] T. S. I. Foundation, "Zonemaster," https://zonemaster.iis.se/en/.

[44] Denic, "Name server predelegation check web interface." [Online]. Available: https://www.denic.de/service/tools/nast/

[45] Cisco, "Cisco prime network registrar 8.2 user guide." [Online]. Available: https://www.cisco.com/c/en/us/td/docs/net_mgmt/prime/network_registrar/8-2/user/guide/CPNR_8_2_User_Guide.html

[46] Miscrosoft, "Troubleshooting DNS Servers." [Online]. Available: https://docs.microsoft.com/en-us/windows-server/networking/dns/troubleshoot/troubleshoot-dns-server#test-a-broken-delegation

[47] S. Hollenbeck, "Extensible Provisioning Protocol (EPP)," RFC 5730, 2009.

[48] "Does Your Domain Have a Registry Lock?" [Online]. Available: https://krebsonsecurity.com/2020/01/does-your-domain-have-a-registry-lock/

[49] "CSC's Research on Election-Related Domains Aligns with Recent FBI and CISA Warning," https://www.circleid.com/posts/20201015-csc-research-election-related-domains-fbi-and-cisa-warning/.

[50] V. Pappas, Z. Xu, S. Lu, D. Massey, A. Terzis, and L. Zhang, "Impact of Configuration Errors on DNS Robustness," in *ACM SIGCOMM*, 2004.

[51] A. J. Kalafut, M. Gupta, C. A. Cole, L. Chen, and N. E. Myers, "An Empirical Study of Orphan DNS Servers in the Internet," in *ACM IMC*, 2010.

[52] A. Phokeer, A. Aina, and D. Johnson, "DNS Lame delegations: A case-study of public reverse DNS records in the African Region," in *AFRICOMM*, 2016.

[53] E. Alowaisheq, S. Tang, Z. Wang, F. Alharbi, X. Liao, and X. Wang, "Zombie Awakening: Stealthy Hijacking of Active Domains through DNS Hosting Referral," in *ACM CCS*, 2020.

[54] T. Vissers, T. Barron, T. Van Goethem, W. Joosen, and N. Nikiforakis, "The Wolf of Name Street: Hijacking Domains Through Their Nameservers," in *ACM CCS*, 2017.

[55] J. G. Kagwe and M. Masinde, "Survey on DNS Configurations, Interdependencies, Resilience and Security for *.Ke Domains," in *the 2nd ACM Symposium on Computing for Development (DEV)*, 2012.

[56] P. Hosein, S. Ramoudith, and K. Mallalieu, "On Internet Resilience in Small Island States," in *6th Int'l Conference on Internet Science*, 2019.

[57] Y. Shi, "Improving E-Government Services Should Start with Domain Names: A Longitudinal Study of Chinese E-Government Domain Names," in *IEEE International Conference on Service Operations and Logistics, and Informatics*, 2007.

[58] M. M. Islam, "Proposed domain name system (DNS) for improved e-government services of Bangladesh," in *12th International Conference on Computers and Information Technology*, 2009.

[59] T. Chung, R. van Rijswijk-Deij, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, and C. Wilson, "Understanding the Role of Registrars in DNSSEC Deployment," in *ACM IMC*, 2017.

[60] V. Visoottiviseth and K. Poonsiri, "The Study of DNSSEC Deployment Status in Thailand," in *IEEE 6th Asian Conference on Defence Technology (ACDT)*, 2019.

[61] S. Singanamalla, E. H. B. Jang, R. Anderson, T. Kohno, and K. Heimerl, "Accept the Risk and Continue: Measuring the Long Tail of Government HTTPS Adoption," in *ACM IMC*, 2020.