

Assessing e-Government DNS Resilience

SIDN Labs Technical Report 2022-01

July 11th, 2022

Raffaele Sommese

University of Twente

r.sommese@utwente.nl

Mattijs Jonker

University of Twente

m.jonker@utwente.nl

Jeroen van der Ham

NCSC-NL / University of Twente

j.vanderham@utwente.nl

Giovane C. M. Moura

SIDN Labs / TU Delft

giovane.moura@sidn.nl

Abstract—Electronic government (e-gov) enables citizens and residents to digitally interact with their government via the Internet. Underpinning these services is the Internet Domain Name Systems (DNS), which maps e-gov domain names to Internet addresses. Structuring DNS with multiple levels of redundancy that can withstand stress events such as denial-of-service (DoS) attacks is a challenging task. While the operator community has established best practices to this end, adopting them all involves expert knowledge and resources. In this work, we obtain and study a list of e-gov domain names used by four countries (The Netherlands, Sweden, Switzerland, and the United States) and measure the DNS structuring of these domains. We show the adoption of best practices, inter-country differences, and provide recommendations for improve DNS service robustness.

Index Terms—DNS, DDoS, E-gov, Resilience, Authoritative DNS

I. INTRODUCTION

Governments increasingly use digital avenues for communication with citizens and residents, further solidifying the Internet as core communications fabric of modern societies. Electronic Governance (e-gov) refers to the set of services governments offer online to their citizens and residents [1]. E-gov has the potential to save costs and provide faster service, easing access to people with disabilities or mobility challenges. The COVID-19 pandemic has exacerbated the payoffs of investments in e-gov, by allowing parts of e-gov services to operate normally despite all restriction measures [1]–[3], such as lockdowns and social distancing.

E-gov depends on the Internet, which in turn relies on the Domain Name System (DNS) [4], [5] as one of its core services. Every web page visit or e-mail that is sent requires DNS resolution. If part of the DNS fails – be it by accident or intentionally as a result of malicious action – domains can become unreachable. One infamous intentional failure resulted from a sizable denial-of-service (DDoS) attack against Dyn [6], a large DNS provider. As a result of this attack, many prominent websites that relied on Dyn – which included Spotify, Netflix, and the New York Times – experienced severe impediment to reachability by their users [7].

As such, e-gov also depends strongly on the DNS, and proper operation of dependent on parts of the DNS is vital to keeping e-gov services accessible. E-gov DNS structuring should therefore be resilient against (partial) failure to avoid service interruption. The DNS supports various levels of

redundancy to become more resilient against events such as DDoS attacks [8], for example through multiple authoritative DNS servers (ADNS) [4], as relying on a single DNS provider creates unnecessary risk. Complementarily, DNS operators can achieve replication through the use of IP anycast [9], [10], which has also proven itself as defensive mechanism. DNS resolvers typically rely on caching [8], which can suppress (temporarily at least) the effects of attacks [11]. Moreover, DDoS filtering and other mitigation techniques can also be deployed to thwart attacks. In spite of this, increasing resilience is not easy. The DNS is prone to many types of configuration errors, which can lead to service unreachability. While best operational practices exist to help increase DNS resilience, some techniques require expert knowledge, operations, and resources – all of which can complicate adopting said practices.

While there have been various studies analyzing DNS infrastructure (e.g., , [11], [12]), there have been few focusing specifically on the DNS infrastructure resilience of -gov domains (e.g., [13]). In this paper, we quantitatively study and evaluate the infrastructure of e-gov DNS, for both web and e-mail services, with regard to DNS and IP-based redundancy. Our goal is to approximate what could happen if these services were to suffer stress events, such as DDoS attacks. We compare the e-gov DNS infrastructure of multiple countries, leveraging the access to the list of e-gov domains that we have. We study three countries in Europe (the Netherlands, Sweden, and Switzerland) as well as the United States in North America. We obtain the lists of e-gov domain names for these countries and use active measurements to evaluate DNS configuration and structuring (§III).

We show that 80% of .gov domain names carry the risk of relying on a *single* DNS provider (§IV). For each of the three continental Europe countries, roughly 40% of the respective domains do so. This risk can be easily remediated by adding additional DNS providers to e-gov domain names. Moreover, the vast majority of domains of all countries are concentrated on a handful providers. The top five providers for each of the four countries are almost exclusively from specifically the country in question.

We also evaluate e-gov domains IP-anycast based replication and DNS caching. We show that most of the continental Europe e-gov domains are not replicated with IP anycast, whereas .gov domains are. The prior should also employ IP

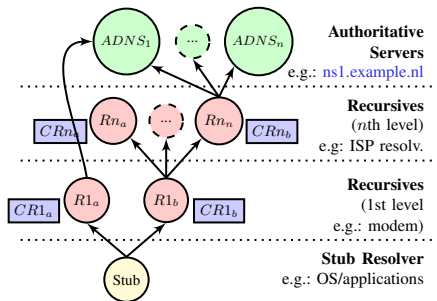


Fig. 1. Relationship between stub resolver (yellow), recursive resolvers (red) with their caches (blue), and authoritative servers (green).

anycast for their DNS to increase their resilience. With regards caching, we show that many e-gov domain DNS infrastructure is configured to not leverage most of the caching features, by setting very low time-to-live (TTL) values for the ADNS servers. Lastly, we analyze the DNS configuration of e-gov domains for incoming mail exchange and investigate which providers domains rely on to this end. We show that Microsoft tops the list for all four countries (§VI).

Overall, we find notable differences among the four evaluated countries with regards to replication of DNS infrastructure. Swiss e-gov domains are lagging behind the other three countries on the use of several best practices.

II. BACKGROUND

In this section, we provide background information on the DNS, different types of DNS servers, and DNS resolution.

Two main types of DNS servers exist. First, *authoritative DNS servers* (ADNS hereafter and in Figure 1), which are those that know the contents of a DNS zone from memory [14]. Second, a recursive resolver (resolver hereafter, R in Figure 1). When a user wants to visit a website, its computer sends a DNS query using its stub resolver (stub in Figure 1) to the recursive, which, on behalf of users, can resolve the requested domain names by querying the multiple ADNS servers needed, through so-called *recursion*. The process is complete when the domain is resolved to its IP addresses: the user can finally visit the requested webpage. Examples of ADNS include the Root DNS servers [15], which are authoritative for the Root zone, and examples of resolvers include Quad[1,8,9] [16]–[18]. In this paper, we analyze the deployment and structuring of ADNS in e-gov.

ADNS have a central role in DNS resolution. Without them, all domains under their respective zones would become *unreachable*. To mitigate this risk, DNS operators can deploy multiple techniques to increase the redundancy and resilience of ADNSes. For example, each DNS zone can use multiple ADNS servers [5] – an example of which is `wikipedia.org`, which uses `ns[1--3].wikimedia.org` as ADNSes. Using IP anycast can further replicate ADNS instances by announcing their respective IP addresses from multiple locations globally. As an example of this consider that the IP address of L-Root, which is one of the 13 ADNSes for the Root Zone, is

	Netherlands	Sweden	Switzerland	United States
TLD	.nl	.se	.ch	.gov
E-gov domains	1309	615	3971	7972
SLD	602	614	3971	7972
FQDN	707	0	0	0

TABLE I
DATASETS FOR WEB DOMAINS (2022-06-08)

announced from 198 locations (June 2022) [15]. Such IP-level replication enables authoritative DNS servers to better cope with DDoS attacks [11], [19]. The last level of replication in DNS is having multiple servers on each anycast location – all behind a load balancer.

On the resolver side, *caching* is the most important technique to offer a safety net for unavailable ADNS [8], [20] (we show caches in Figure 1 as CR). Whenever a resolver queries a ADNS, it keeps in a memory cache the responses. ADNS operators specify the maximum caching time using the time-to-live (TTL) value of the response [4]. Caching not only protects users from ADNS DDoS, but also improve response times by having cache hits. However, the safety net holds only as long as records are cached.

The DNS is used for more than IP address resolution. Enabling domain-destined (incoming) e-mail relies on a specific DNS record: the Mail Exchanger (MX) record. MX records are provided in a label format (*e.g.*, `google.com` has `smtp.google.com` as MX label), which must be resolved to by the sending Mail Transfer Agent (MTA) to determine the location (IP address) of the receiving MTA. As e-mail service can be provisioned by a third party, the authoritative DNS infrastructures for MX label and MX address resolution are not necessary the same.

III. DATASETS AND MEASUREMENTS

A. Datasets

Table I shows the datasets for web domains we employ in this paper. We obtain e-gov domain names from the Netherlands, Sweden, Switzerland, and the United States. By and large, the e-gov domains we consider are second-level domains (SLD) such as `cdc.gov`. Thanks to our collaboration with the Netherlands’ National Cyber Security Center (NCSC-NL) in this research, we were also able to obtain fully-qualified domain names (FQDN) associated with e-identity services in the Netherlands. These are typically fully-qualified domain names such as `login.smalltown.example.nl`. We treat them as a different category when we compare them with other countries. We obtained Sweden’s e-gov domain list from the Swedish Internet Foundation (IIS), which operates `.se`. The Swiss e-gov domains were provided to us by SWITCH, the `.ch` registry. For the United States’ e-gov domains, we analyzed the complete list of `.gov` names, provided in a public dataset.¹

E-mail E-gov domains may also be used for e-mail (*e.g.*, `info@nsf.gov`), which involves MX record resolution

¹<https://home.dotgov.gov/data/>

(see §II), and may involve an external e-mail provider. We evaluate external dependencies for e-gov domains. For the Netherlands, we obtained a specific list of domains that are used for e-mail. For the other countries, we target obtained e-gov names with MX queries.

B. Measurements

We instrumented our own DNS measurements for this study. We carried out these measurements from a single vantage point on 2022-06-07. This vantage point is provisioned on the Netherlands National Research and Education Network. Our measurement sequence is as follows (Figure 2). For each domain d , we measure:

- 1) The respective set of ADNS servers, NS_d , from both parent and child authoritative servers, as defined by d 's NS records in DNS (see 1 – 2 – 3 – 4 in Figure 2)
- 2) For each $ns_d \in NS_d$, we then query for the IPv4 and IPv6 address records, as defined by ns_d 's A and AAAA records in DNS (see 5 – 6 in Figure 2)
- 3) We then target each ns_d IP address measured with two queries related to d as follows (see also 7–8 in Figure 2):
 - a) An A query to determine if ns_d is not misconfigured for the respective d (e.g., lame delegation)
 - b) An MX query to obtain the mail exchanger records (MX records) of d , MX_d ²
- 4) We then measure the ADNS records, NS_{mx_d} , of each $mx_d \in MX_d$, query for the IP addresses of all $ns_{mx_d} \in NS_{mx_d}$, and use the measured IP addresses to perform a lame delegation check on ns_{mx_d} for mx_d

For each ADNS IP address learned in the previous step, we run additional measurements to determine if they are IP anycast. To perform this task, we utilize iGreedy [21]. iGreedy detects anycast prefixes using the great circle distance methodology (GCD). Running round-trip-time (RTT) measurements from geographically distributed VPs makes it possible to detect anycast instances by using speed of light constraint violations. We use 500 RIPE Atlas probes as globally distributed vantage points for the measurements. The probes we select are all a minimum distance of 100km apart.

We chose to measure IP anycast use ourselves rather than use on publicly available data. While a public anycast census exists [22], it lacks IPv6 information. Regarding geolocation, the iGreedy measurement mechanics offer the means to determine this for anycast IP addresses. For addresses classified as unicast we rely on IP2Location for geolocation.

C. Limitations

Our study knows several limitations. First, we perform DNS resolution from a single vantage point in the Netherlands, which may introduce bias if the targeted servers filter or change the responses based on the requesting IP. Second, our anycast census leverages ICMP reachability of the targeted IPs, which as discussed in [21]–[24] can lead to a lower-bound estimation of anycast deployment. Finally, our analysis

²For Netherlands' e-gov names we only target those known to use e-mail.

	NL	SE	CH	GOV
E-gov domains	1309	615	3971	7972
SLD	602	614	3971	7972
Responsive	601	609	3546	7911
single provider(v4/v6)	268/331	249/254	1531/1923	6564/4455
multi-provider(v4/v6)	333/266	360/254	2013/344	1306/578

TABLE II
RESPONSIVE DOMAINS (2022-06-08)

is scoped to only four TLDs, for which we were able to obtain lists of e-gov domains.

IV. SINGLE DEPENDENCIES

In this section, we focus on e-gov domains' singular dependency on providers and infrastructure. Given that DNS is highly distributed, we analyze the dependency of individual elements of the infrastructure as possible cause of unreachability of the e-gov domains in case of failure.

A. ADNS providers dependency

We start by analyzing the number of DNS providers of each e-gov domain. For each domain, d , we first measure the ADNS servers (§III-B). In the case of `example.nl` (Figure 3), that would be two ADNS *server* names: `a.example.nl` and `b.example.com`. Then we resolve the ADNS server's IP addresses, using A and AAAA type queries [25]. For example, in Figure 3, `a.example.nl` has 192.168.1.1 as IP address. For each IP address, we then look up its Autonomous System (AS) [26] number. Then we compute the number of unique ASes for d (for IPv4 and IPv6, separately). In our example, `example.nl` has two ADNS providers: AS1234 and AS456.

Table II shows the number of analyzed e-gov domains. The number of responsive domains is slightly smaller than the number of *actual* e-gov SLDs. There are multiple reasons for this. Some domains have *delegation problems* [27]: some list wrong ADNS servers (servers that are not responsive or are not authoritative for the domain name). For example, `daviscountyutah.gov` lists 168.180.200.18 as IP address of one of its ADNS (`dc-dns.daviscountyutah.gov`). However, this IP does not respond to DNS queries. Other domains, such as `hudsoncountynj.gov`, although listed as a `.gov` domain names, have already been removed from the zone, so they do not exist.

For each responsive domain on Table II, we compute the number of ADNS providers (as measured by their AS numbers). Figure 5 shows the CDF of DNS providers for the responsive e-gov domains. For the ccTLDs (`.nl`, `.se`, and `.ch`), we notice that roughly 40% of the e-gov domains have a single ADNS provider. For `.gov`, however, the majority of domains (80%+) have a single ADNS provider.

1) *ADNS provider consolidation and centralization*: Next, we focus on the side-effects of using third-party DNS providers: shared infrastructure. For each country, we compute the number of e-gov domains that each ADNS provider has. Figure 4 shows the results. We can see that regardless of the zone, a handful of DNS providers exclusively operate the

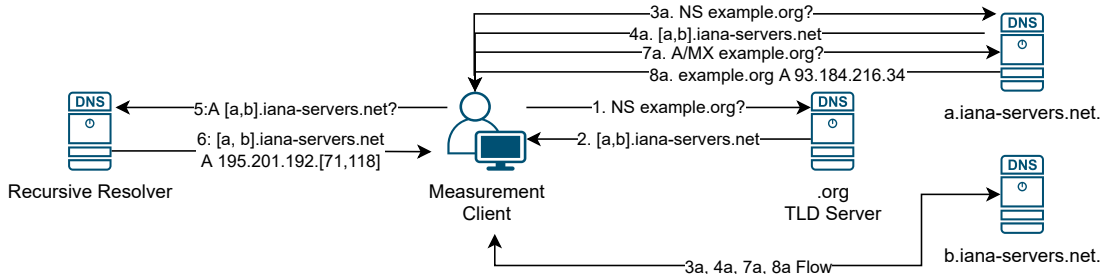


Fig. 2. Measurements step explained: Our software first query the TLD server (PARENT) and following all the CHILDREN servers.




Fig. 3. Relationship between domain names and DNS records

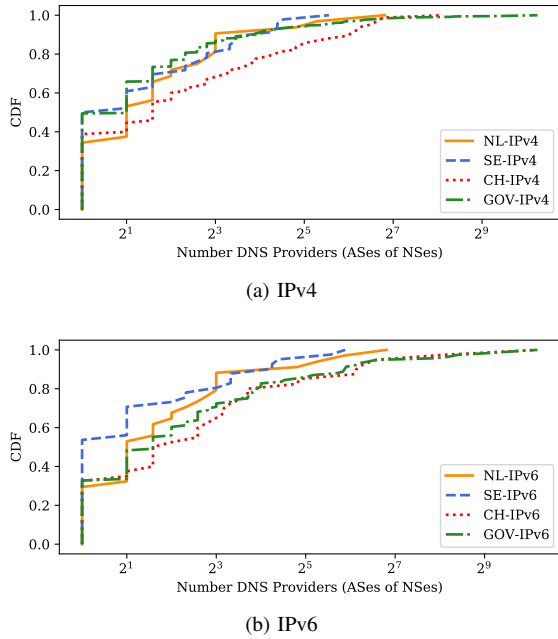



Fig. 4. ADNS provider domains concentration


majority of the domains. Table III shows, per DNS zone, the top ADNS providers and the number of e-gov domains hosted.

We also see that *local* DNS providers provide service to most of the domains, *i.e.*, DNS providers from the countries' e-gov – the exception being Microsoft showing up as #3 for Sweden (we manually verified these domains and they use Microsoft Azure DNS servers, such as ns3-02.azure-dns.org).

These results show that although there has been a growing consolidation and centralization of DNS infrastructure over the last years in the hands of large US-based companies [12], [28], this has not been the case for the continental European e-gov domains we study.



(a) IPv4



(b) IPv6

Fig. 5. Number of ASes (DNS providers) for e-gov domains

Implications: although relatively rare, large DNS providers can have (partial) failures – as in the case of Dyn and AWS [29]. In case of a massive DDoS attack on the provider, the associated e-gov domains may experience serious reachability issues (clients will not be able to resolve them). While having a single DNS provider may simplify the configuration, it places the reliability of domains into a single provider. Given that DNS providers share infrastructure among all their DNS zones, even attacks on another domains can bring down e-gov domains, due to collateral damage. As such, as also pointed out by Allman [30], it is better to use multiple ADNS providers and even operate one of them in-house. Our contribution is to measure it for e-gov domains from multiple countries – and show that more .gov e-gov domains depend on a single ADNS provider than the other ccTLDs, which share similar rates.

B. ADNS servers dependency

To avoid single points of failure, the original DNS RFC (RFC1034 [4]) requires that domain names have at least two ADNS servers, as in Figure 3. We evaluate this requirement

	NL		SE		CH		GOV	
	ASN	e-gov	ASN	e-gov	ASN	e-gov	ASN	e-gov
#1	20857 - Transip (NL)	112	39570 - Loopia (SE)	47	29222 - Infomaniak (CH)	278	44273 - GoDaddy (US)	1215
#2	48635 - CLDIN (NL)	39	1257 - Tele2 (SE)	23	3303 - Swisscomm (CH)	115	13335 -Cloudflare (US)	909
#3	12315 - QSP (NL)	28	8068 - Microsoft (US)	21	35206 - Novatrend (CH)	100	16509 - Amazon (US)	676
#4	29311 - Solvinity (NL)	8	1729 - Telia (SE)	21	9108 -Abraxas (CH)	97	21342 - Akamai (US)	334
#5	48037 - SSC-ICT (NL)	8	3301 - Telia (SE)	19	21069 - Metanet (CH)	91	16552 - Tiggee (US)	316

TABLE III
TOP ADNS PROVIDERS CONCENTRATION FOR SINGLE ADNS E-GOV DOMAINS (IPv4)




Fig. 6. Number of ADNS servers (NS records) for e-gov domains

in terms of namespace – *i.e.*, as two different ADNS names (NS records in Figure 3) but also as different network prefixes. If two ADNS servers share the same prefix, they are announced from the same location and, therefore, share the same infrastructure and are not topologically diverse.


We start with the ADNS server name analysis. Figure 6 shows the CDF of the domain names and their respective numbers of NS records. We see that the vast majority of e-gov domains have at least two ADNS servers, conforming to RFC1034. We found one .ch that had only one NS record at the child delegation (§III), but two at the parent .ch delegation. This is caused by well-known parent/child inconsistency [27].

We also identify 37 .gov domains with a single ADNS name in their child delegation. Out of these, 32 had more than NS in their parent .gov authoritative server – but 5 did not. The .gov stipulates that their domains must have two ADNS servers³). However, these six domains⁴ do not conform to it. As such, these 37 violates what RFC1034 stipulates, and six violate .gov policy. We notified the .gov registry and registrar of this.


While most e-gov domains have at least two ADNS servers (two different NS records), we now determine if this redundancy is also found on their associated IP prefixes. For each IP address, we retrieve their BGP prefix using CAIDA Prefix-to-AS mapping [31] and compute the number of prefixes each e-gov domain has. Figure 7 shows the results. We see that Switzerland’s .ch e-gov domains lead the number of domains with a single BGP prefix (also shown in Table IV) – roughly one-third of its e-gov domains ADNS servers on the same network prefix. For IPv6, it is even worse: roughly 40% of

³See .gov requirements at: <https://home.dotgov.gov/help/#what-are-the-name-server-requirements-for-gov-domains>

⁴these six are: theftaz.gov, ncrealid.gov, bardstownky.gov, sjcpa.gov, cityofdelafieldwi.gov., villageofpewaukee.gov



(a) IPv4



(b) IPv6

Fig. 7. Number of distinct BGP prefixes that announce ADNS IP addresses

	NL	SE	CH	GOV
Responsive	601	609	3546	7911
Single prefix(v4/v6)	125/341	127/203	1078/1748	1241/885
Anycast(v4/v6)	125/125	77/77	87/81	4425/3643

TABLE IV
E-GOV DOMAINS, PREFIXES AND ANYCAST USAGE

the domains do not support DNS over IPv6, and another 40% are announced from a single prefix.

Implications: RFC1034 states that ADNS servers for the same DNS zone should be placed in topologically distinct networks. We have seen that many e-gov domains, for all zones, depend on ADNS servers located in the same location. This creates an unnecessary risk in case of failures or attacks. As such, we recommend these operators configure ADNS servers in other distinct networks. Note that simply having different prefixes does not guarantee topological diversity [30], but having the same prefix implies lack of topological diversity.

C. TLD dependency

Next, we investigate what top-level domains (TLD) the ADNS servers of e-gov domains depend upon. For example, in Figure 3, we see that example.nl’s NS records end in .nl and .gov, so it depends on two TLDs. While TLDs failures are




Fig. 8. Number of TLDs used by e-gov domains ADNS

	NL	SE	CH	GOV
1	170 (.nl)	483 (.se)	609 (.ch)	2507 (.com)
2	69 (.net)	100 (.net)	190 (.com)	1541 (.net)
3	26 (.com)	82 (.com)	150 (.net)	894 (.gov)
4	12 (.eu)	14 (.info)	19 (.org)	485 (.org)
5	4 (.be)	8 (.org)	12 (.de)	302 (.us)

TABLE V
MOST USED TLD BY E-GOV ADNS SEVERs.

unlikely (just as large cloud provider failures), it is important to prepare for them and avoid that a potential TLD failure leads to e-gov domain unreachability.


For each e-gov domain, we compute the number of TLDs on which they depend, by analyzing the name of its ADNS servers. We then generate the CDFs per country, shown in Figure 8. We see that Swiss e-gov domains are heavily concentrated in one TLD. The United States’ .gov e-gov ADNS servers are also heavily concentrated, followed by Sweden and the Netherlands.

Table V shows the top 5 TLDs for each country. To calculate this, we first generate a list of all ADNS servers for e-gov domains per country. Then, we extract their TLDs and count and rank them. We see that cultural affinity seems to play a role in these results. The three countries from continental Europe use mostly their own countries’ ccTLD, followed by either .net or .com (which are present in all 4 countries’ e-gov domains). The US’s .gov most rely on .com domains, given it is a TLD operated in the US (as is all 5 in the .gov list) and where most cloud providers register DNS names.


Implications: e-gov domains could benefit from extra resilience by having ADNS servers with FQDNs under a more diverse set of TLDs. This can protect such domains from failures in TLDs. Although this may be unlikely, these extra measures do not add much extra complexity and provide extra resilience. To illustrate this in practice, consider the domain `digid.nl`, which provides Dutch citizens with e-identity services to access their e-gov services. This domain uses .com, .nl, .org and .eu as TLDs in its ADNS servers.

V. ANYCAST AND CACHING IN E-GOV

In the previous section we focused on analyzing e-gov domain dependency on various parts of the Internet infrastructure. In this section, we focus on if and how e-gov domains rely on two particular techniques to improve resilience: the use of IP anycast and then of DNS caching.



(a) IPv4



(b) IPv6

Fig. 9. Anycast Adoption by e-gov domains

A. Anycast adoption

IP anycast is one of the cornerstones of DNS resilience. As such, DNS operators should deploy anycast to have more robust ADNS services [19]. For this reason, we quantify anycast adoption among e-gov domains.

Figure 9 shows the CDF of e-gov domains with regards anycast adoption. We see major differences between the countries under study. Around 58% of .gov domains have one or more anycast ADNS servers, whereas very few Swiss e-gov domains do. The Netherlands and Sweden score in between; approximately 15–20% of their e-gov domains have at least one ADNS that is anycast. The reason for this, we believe, has to do with the ADNS providers. .gov is mostly served by large ADNS providers (Table III) whereas the other ccTLDs are mostly served by local companies, which may not deploy anycast or may charge an additional fee for this service.

Implications: IP anycast is widely deployed to improve DNS resilience. We see that most of the continental European e-gov domains under consideration do not support anycast, while the majority of the US .gov domains do. We hope that the European and Swiss e-gov domains will in the future start using anycast.

B. Caching

DNS resolvers heavily deploy caching of DNS responses to improve response times to clients (Figure 1). It is by far the most efficient method to cut response times [20] and it can even suppress the effects of DDoS attacks [8], as clients can still resolve domain names thanks to cache hits, even when ADNS servers are unreachable.

While DNS caching is performed by resolvers, it is ADNS that controls how long records should stay in DNS resolver

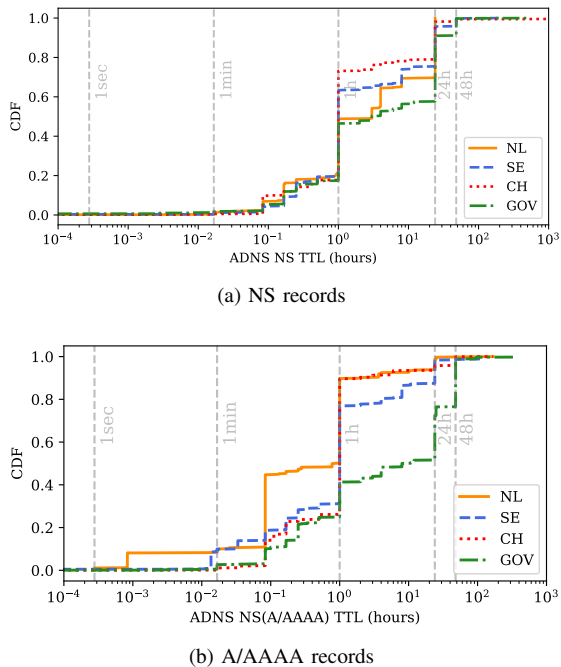


Fig. 10. TTLs usage in ADNS

	NL	SE	CH	GOV
NS TTL				
1st quartile	3600	3600	3600	3600
Median	10800	3600	3600	10800
3rd quartile	86400	40001	10800	86400
A/AAAA TTL				
1st quartile	300	900	1800	3600
Median	3047	3600	3600	28800
3rd quartile	3600	3600	3600	90000

TABLE VI
TTL DISTRIBUTION (s)

cache – by setting a time-to-live (TTL) value on each DNS record under their DNS zones. While TTL values range from 0s to years, in practice most records fall between 10min and 24h [20]. It is suggested to configure ADNS NS records to have a TTL of at least a few hours [19], [20].

We next analyze the TTLs of both ADNS NS records and their respective IP address records (e.g., A and AAAA). Figure 10 shows the results for the zones we evaluate. For NS records, we see that most records fall between 1 and 24h and many are equal to 1h, which is considered *short* for a NS record [19], [20]. For A/AAAA records, we see that most fall within an interval of up to one hour.

Table VI shows the TTL distributions. We see that the median TTL for Sweden and Switzerland is rather low (1h) for the NS records. For A/AAAA records, we see that most domains have a 1h TTL, which is considered reasonable.

Implications: Caching is the last line of defense for e-gov domains: if all ADNS servers for a domain are down, a client may still be able to resolve and reach the e-gov website if its resolver has the domain in question in cache. We see that the NS TTLs from Sweden and Switzerland (1h median) may be too low. The Netherlands and US have a 3h median. These

zones could benefit from longer caching if their TTLs are increased. However, such a configuration change must take into account if the DDoS protection services do not depend on DNS redirection, which benefits from shorter TTLs.

VI. EXTERNAL MAIL DEPENDENCY

Recall from §II that labels in MX records must be resolved to determine the location of the receiving mail server, which can involve “external” ADNS infrastructure. As DNS infrastructure involved in MX resolution should also be resilient, we evaluate dependencies for the obtained e-gov domains.

Table VI shows the dependencies. Out-of-zone means that MX label resolution involves external ADNS infrastructure. Same zone means that the MX label is in the same zone as the e-gov domain (e.g., `mx-west.fbi.gov` is in the zone of `fbi.gov`). Mixed means a combination, which can occur in case an e-gov domain defines multiple MX records. We observe that e-gov domains heavily rely on external ADNS infrastructure. The smallest percentage of same zone MX labels is seen for `.gov` (12.6%). We also see a small number of cases where two records are combined.

For e-gov domains relying on third-party mail providers, we further investigate the mail provider, identified by SLD. As shown in §VI, Microsoft Outlook services prominently services e-gov domains. We also observe several in-country mail providers for `.nl` and `.ch`. For example, for `.nl`, `ssonet.nl` is a large IT provider of the Netherlands Government.

Considering anycast of MX ADNS infrastructure, we observe that a significant percentage (87.5%) of third-party mail providers use anycast for their ADNS servers. We find a comparable percentage for `.se` and `.ch`. For `.gov` names, we see lower (62% of 5944 FQDNs) anycast adoption.

To study the resilience of out-of-zone dependencies in terms of *network diversity*, we perform a case study for `.nl` providers. Among the MX labels for `.nl` e-gov names, we identify 330 unique FQDN (i.e., MX labels). Our measurement data for these labels shows at least two NS records and two v4 ADNS servers for all labels, but only two-thirds with at least two v6 ADNS servers. Of the resolved ADNS infrastructure addresses, 66% are hosted in a single ASN for IPv4, and 72% for IPv6. All the v4 authoritative nameservers responded and only 2% of the v6 authoritative nameserver did not.

Implications: Third-party e-mail providers on which e-gov names depend offer, for the most, resilient ADNS infrastructure, hardening the additional resolution step for MX labels.

VII. DISCUSSION AND RECOMMENDATIONS

A robust, redundant, and properly configured DNS is crucial for e-gov services to be delivered to citizens and residents. We compare four countries with regards their e-gov DNS structuring. Our results show that there is plenty of room for improvement, which we cover next.

First and foremost, we show that there is much dependency on single DNS providers, for all countries under study (§IV-A). The e-gov domains should add at least a second DNS provider, which could protect them against failure and attacks that can

MX Provider	#.nl Domains	%.nl Domains	MX Provider	#.se Domains	%.se Domains
outlook.com	164	(39.0%)	outlook.com	205	(37.5%)
ezorg.nl	46	(11.0%)	mailanyone.net	69	(12.6%)
ssonet.nl	17	(4.0%)	mx25.net	52	(9.5%)
barracudanetworks.com	13	(3.1%)	staysecuregroup.com	38	(6.9%)
minvenj.nl	12	(2.9%)	staysecuregroup.net	38	(6.9%)
MX Provider	#.ch Domains	%.ch Domains	MX Provider	#.gov Domains	%.gov Domains
outlook.com	425	(22.1%)	outlook.com	2243	(41.4%)
infomaniak.ch	129	(6.7%)	google.com	532	(9.8%)
abxsec.com	120	(6.2%)	barracudanetworks.com	495	(9.1%)
tophost.ch	90	(4.7%)	pphosted.com	161	(3.0%)
ag.ch	78	(4.1%)	mimecast.com	157	(2.9%)

TABLE VII
TOP 5 THIRD-PARTY E-MAIL PROVIDER PER COUNTRY

TLD	Mail Domains	Same zone	Mixed	Out-of-zone
.gov	5797	733 (12.6%)	121 (2.1%)	4943 (85.3%)
.ch	2126	302 (14.2%)	10 (0.5%)	1841 (85.3%)
.se	544	113 (20.8%)	5 (0.9%)	426 (78.3%)
.nl	508	102 (20.1%)	5 (1%)	401 (78.9%)

TABLE VIII
#DOMAINS RELYING ON SAME ZONE, MIXED OR OUT-OF-ZONE DNS INFRASTRUCTURE FOR MX LABEL RESOLUTION.

occur within individual providers (an event seen multiple times in the past).

Secondly, we observe that many e-gov domains have ADNS infrastructure in the same networks (§IV-B) – violating recommendations from the original DNS RFCs. We recommend e-gov domains to adhere to these recommendations.

Third, we found that for the evaluated countries in continental Europe, DNS service is largely provided by local providers, and not by the large US-based cloud and DNS providers (§IV-A). We can only speculate that this may be due to historical reasons – the large US-based cloud services are relatively new compared to most e-gov domains. For e-gov e-mail, however, it is completely different: Microsoft dominates the e-gov market in all countries – which could be due to the usage of Outlook’s cloud-based e-mail services.

Our final recommendation is for operators to carefully set the TTL values of their DNS records, so they can leverage the benefits of caching in DNS during stress events §V. This requires only a single parameter change. Similarly, we also recommend that countries deploy more IP anycast on their ADNS servers. We show that despite having the highest GDP per capita, Switzerland lags behind in terms of anycast adoption. We will present our findings to the respective countries TLD operators.

VIII. RELATED WORK

DNS and e-gov: The closest research work to ours is by Houser et al. [13], who also investigate e-gov DNS. They look into web domains while we look into both e-mail and e-mail DNS infrastructure. They cover government domains of 193 countries and use, like us, active measurements to measure ADNS infrastructure. We, however, differ in several ways:

first, the input domains: we obtain a list of e-gov domain names either publicly (.gov and .se) or privately (.se and .ch) – so we have a complete view of these zones. Houser et al., however, focus on inferring domain names using a combined set of methods. While their coverage is larger in terms of TLDs, it may miss e-gov domains. Ours, in turn, cover only four countries but with a complete view of their domains. They analyze zone inconsistencies and delegation errors. We focus on e-gov DNS structuring from a stress event angle.

DNS resilience: standardization efforts and ample research exists for DNS resilience and redundancy. RFC9199 [19] summarizes six considerations for large ADNS operators – including using IP anycast on every single ADNS [11], [32], optimizing routing is then more important than adding extra locations on anycast networks [33], and considering long TTL values to leverage the benefits of resolver’s caching [8], [20]. *DNS consolidation and centralization:* Allman previously studied ADNS replication and shared infrastructure for .com and .net [30]. Similar to us, they recommend using multiple ADNS providers and to deploy topologic diverse ADNSes. They found that 28% of the second-level domains do not meet the multiple networks requirement for ADNS diversity. Another study investigated the now-defunct top 100k Alexa [34] domains and their ADNS infrastructure [12], showing that 89% of the domains rely on managed DNS providers, and 28% use a single DNS provider. The authors also showed how 3 DNS providers host 40% of the 100k Alexa websites. Centralization on the resolver market has also been quantified from the Netherlands .nl ccTLD [28]. The authors found that one-third of the queries to .nl domain names originate from five large cloud/content providers.

IX. CONCLUSIONS

E-gov has become an essential part of government. As a core Internet protocol, the DNS underpins reachability of e-gov services. In this paper we evaluated DNS structuring for e-gov services (web and e-mail) for four countries. Our results show that many e-gov domains are not following the current recommendations for operation of large DNS providers, regardless of country. While e-gov domains may

operate without hiccups, it is not free of risks, as a motivated attacker could stress specific DNS infrastructures to deteriorate the reachability of many e-gov domains. We hope our findings prompt the responsible operators to improve the redundancy and resilience of e-gov DNS.

Acknowledgements This work was supported by: the DINO project, contracted by the Netherlands’ National Cyber Security Center (NCSC-NL); the EU H2020 CONCORDIA project (830927); and the joint US Department of Homeland Security and Dutch Research Council DHS-NWO MADDVIPR project (628.001.031/FA8750-19-2-0004).

REFERENCES

- [1] M. Gotze, S. Matic, C. Iordanou, G. Smaragdakis, and N. Laoutaris, “Measuring Web Cookies in Governmental Websites,” in *14th ACM Web Science Conference 2022*, ser. WebSci ’22. New York, NY, USA: Association for Computing Machinery, 2022, p. 44–54. [Online]. Available: <https://doi.org/10.1145/3501247.3531545>
- [2] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis, “A Year in Lockdown: How the Waves of COVID-19 Impact Internet Traffic,” *Commun. ACM*, vol. 64, no. 7, p. 101–108, jun 2021.
- [3] O. for Economic Co-operation and Development, “Responding to COVID-19: The rules of good governance apply now more than ever,” <https://www.oecd.org/governance/public-governance-responses-to-covid19/>, Jun. 2022.
- [4] P. Mockapetris, “Domain names - concepts and facilities,” IETF, RFC 1034, Nov. 1987. [Online]. Available: <http://tools.ietf.org/rfc/rfc1034.txt>
- [5] —, “Domain names - implementation and specification,” IETF, RFC 1035, Nov. 1987. [Online]. Available: <http://tools.ietf.org/rfc/rfc1035.txt>
- [6] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, “Understanding the Mirai botnet,” in *Proceedings of the 26th USENIX Security Symposium*. Vancouver, BC, Canada: USENIX, Aug. 2017, pp. 1093–1110. [Online]. Available: <https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-antonakakis.pdf>
- [7] N. Perlroth, “Hackers used new weapons to disrupt major websites across U.S.” *New York Times*, p. A1, Oct. 22 2016. [Online]. Available: <http://www.nytimes.com/2016/10/22/business/internet-problems-attack.html>
- [8] G. C. M. Moura, J. Heidemann, M. Müller, R. de O. Schmidt, and M. Davids, “When the dike breaks: Dissecting DNS defenses during DDoS,” in *Proceedings of the ACM Internet Measurement Conference*. Boston, MA, USA: ACM, Oct. 2018, pp. 8–21.
- [9] C. Partridge, T. Mendez, and W. Milliken, “Host Anycasting Service,” IETF, RFC 1546, Nov. 1993. [Online]. Available: <http://tools.ietf.org/rfc/rfc1546.txt>
- [10] D. McPherson, D. Oran, D. Thaler, and E. Osterweil, “Architectural Considerations of IP Anycast,” IETF, RFC 7094, Jan. 2014. [Online]. Available: <http://tools.ietf.org/rfc/rfc7094.txt>
- [11] G. C. M. Moura, R. de O. Schmidt, J. Heidemann, W. B. de Vries, M. Müller, L. Wei, and C. Hesselman, “Anycast vs. DDoS: Evaluating the November 2015 root DNS event,” in *Proceedings of the ACM Internet Measurement Conference*. Santa Monica, California, USA: ACM, Nov. 2016, pp. 255–270.
- [12] 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 *Proceedings of the ACM Internet Measurement Conference*, ser. IMC ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 634–647.
- [13] R. Houser, S. Hao, C. Cotton, and H. Wang, “A Comprehensive, Longitudinal Study of Government DNS Deployment at Global Scale,” in *Proceedings of the 52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. IEEE/IFIP, jun 2022.
- [14] P. Hoffman, A. Sullivan, and K. Fujiwara, “DNS Terminology,” IETF, RFC 8499, Jan. 2019. [Online]. Available: <http://tools.ietf.org/rfc/rfc8499.txt>
- [15] Root Server Operators, “Root DNS,” 2022, <http://root-servers.org/>.
- [16] 1.1.1.1, “The Internet’s Fastest, Privacy-First DNS Resolver,” <https://1.1.1.1/>, Jun. 2022. [Online]. Available: <https://1.1.1.1/>
- [17] Google, “Google Public DNS,” Jun. 2022. [Online]. Available: <https://developers.google.com/speed/public-dns/>
- [18] Quad9, “A public and free DNS service for a better security and privacy,” <https://quad9.net>, Jun. 2022.
- [19] G. Moura, W. Hardaker, J. Heidemann, and M. Davids, “Considerations for Large Authoritative DNS Server Operators,” IETF, RFC 9199, Mar. 2022. [Online]. Available: <http://tools.ietf.org/rfc/rfc9199.txt>
- [20] G. C. M. Moura, J. Heidemann, R. de O. Schmidt, and W. Hardaker, “Cache me if you can: Effects of DNS Time-to-Live,” in *Proceedings of the ACM Internet Measurement Conference*. Amsterdam, the Netherlands: ACM, Oct. 2019, pp. 101–115.
- [21] D. Cicalese and D. Rossi, “A Longitudinal Study of IP Anycast,” *SIGCOMM Comput. Commun. Rev.*, vol. 48, no. 1, pp. 10–18, Apr. 2018.
- [22] R. Sommesse, L. Bertholdo, G. Akiwate, M. Jonker, van Rijswijk-Deij, Roland, A. Dainotti, K. Claffy, and A. Sperotto, “MANycast2—using anycast to measure anycast,” in *Proceedings of the ACM Internet Measurement Conference*. Pittsburgh, PA, USA: ACM, Oct. 2020.
- [23] J. Heidemann, Y. Pradkin, R. Govindan, C. Papadopoulos, G. Bartlett, and J. Bannister, “Census and Survey of the Visible Internet,” in *Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC ’08. New York, NY, USA: ACM, 2008.
- [24] R. Sommesse, G. Akiwate, M. Jonker, G. Moura, M. Davids, R. van Rijswijk - Deij, G. Voelker, S. Savage, K. Claffy, and A. Sperotto, “Characterization of anycast adoption in the dns authoritative infrastructure,” in *5th Network Traffic Measurement and Analysis Conference, TMA 2021*. IFIP, Sep. 2021.
- [25] S. Thomson, C. Huitema, V. Ksinant, and M. Souissi, “DNS Extensions to Support IP Version 6,” IETF, RFC 3596, Oct. 2003. [Online]. Available: <http://tools.ietf.org/rfc/rfc3596.txt>
- [26] Y. Rekhter, T. Li, and S. Hares, “A Border Gateway Protocol 4 (BGP-4),” IETF, RFC 4271, Jan. 2006. [Online]. Available: <http://tools.ietf.org/rfc/rfc4271.txt>
- [27] R. Sommesse, 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 *Passive and Active Measurement*. Cham: Springer International Publishing, 2020, pp. 175–189.
- [28] G. C. M. Moura, S. Castro, W. Hardaker, M. Wullink, and C. Hesselman, “Clouding up the internet: How centralized is dns traffic becoming?” in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC ’20. New York, NY, USA: Association for Computing Machinery, 2020, p. 42–49.
- [29] C. Williams, “Bezos DDoS’d: Amazon Web Services’ DNS systems knackered by hours-long cyber-attack,” https://www.theregister.co.uk/2019/10/22/aws_dns_ddos/, 10 2019.
- [30] M. Allman, “Comments on DNS Robustness,” in *Proceedings of the Internet Measurement Conference 2018*, ser. IMC ’18. New York, NY, USA: Association for Computing Machinery, 2018, p. 84–90. [Online]. Available: <https://doi.org/10.1145/3278532.3278541>
- [31] CAIDA, “Routeviews Prefix to AS mappings Dataset for IPv4 and IPv6,” 2020. [Online]. Available: <http://www.caida.org/data/routing/routeviews-prefix2as.xml>
- [32] M. Müller, G. C. M. Moura, R. de O. Schmidt, and J. Heidemann, “Recursives in the wild: Engineering authoritative DNS servers,” in *Proceedings of the ACM Internet Measurement Conference*. London, UK: ACM, 2017, pp. 489–495. [Online]. Available: <http://www.isi.edu/%7ejohnh/PAPERS/Mueller17b.html>
- [33] R. d. O. Schmidt, J. Heidemann, and J. H. Kuipers, “Anycast latency: How many sites are enough?” in *Proceedings of the Passive and Active Measurement Workshop*. Sydney, Australia: Springer, Mar. 2017, pp. 188–200. [Online]. Available: <http://www.isi.edu/%7ejohnh/PAPERS/Schmidt17a.html>
- [34] Alexa, “Alexa: Keyword Research, Competitive Analysis & Website Ranking,” Feb. 2022. [Online]. Available: <https://www.alexa.com/>