

How Ready Is DNS for an IPv6-Only World?

Florian Streibelt¹, Patrick Sattler², Franziska Lichtblau¹, Carlos H. Gañán³,
Anja Feldmann¹, Oliver Gasser¹, and Tobias Fiebig¹

¹ Max Planck Institute for Informatics

{fstreibelt,rhalina,anja,oliver.gasser,tfiebig}@mpi-inf.mpg.de

² TU München sattler@net.in.tum.de

³ TU Delft c.hernandezganan@tudelft.nl

Abstract. DNS is one of the core building blocks of the Internet. In this paper, we investigate DNS resolution in a strict IPv6-only scenario and find that a substantial fraction of zones cannot be resolved. We point out, that the presence of an **AAAA** resource record for a zone’s nameserver does not necessarily imply that it is resolvable in an IPv6-only environment since the full DNS delegation chain must resolve via IPv6 as well. Hence, in an IPv6-only setting zones may experience an effect similar to what is commonly referred to as lame delegation.

Our longitudinal study shows that the continuing centralization of the Internet has a large impact on IPv6 readiness, i.e., a small number of large DNS providers has, and still can, influence IPv6 readiness for a large number of zones. A single operator that enabled IPv6 DNS resolution—by adding IPv6 glue records—was responsible for around 15.5% of all non-resolving zones in our dataset until January 2017. Even today, 10% of DNS operators are responsible for more than 97.5% of all zones that do not resolve using IPv6.

1 Introduction

With the recent exhaustion of the IPv4 address space, the question of IPv6 adoption is gaining importance. More end-users are getting IPv6 prefixes from their ISPs, more websites are reachable via IPv6, hosting companies start billing for IPv4 connectivity or give discounts for IPv6-only hosting and IoT devices further push IPv6 deployment. Yet, one of the main entry-points for Internet services—the DNS—is suffering from a lack of pervasive IPv6 readiness. While protocols such as Happy Eyeballs [41, 45] help to hide IPv6 problems, they complicate detection and debugging of IPv6 issues. Indeed, the threat of DNS name space fragmentation due to insufficient IPv6 support was already predicted in RFC3901, over 18 years ago [18]. Hence, in this paper, we measure the current state of IPv6 resolvability in an IPv6-only scenario.

In Figure 1 we show two common misconfigurations, which prevent DNS resolution over IPv6 and lead to an effect similar to what is commonly called lame delegation. Note, that RFC8499 [26] defines lame delegation as incorrect NS entries or nameservers *not responding properly*. While the observed behaviour might look the same, the underlying misconfiguration, e.g., missing AAAA or

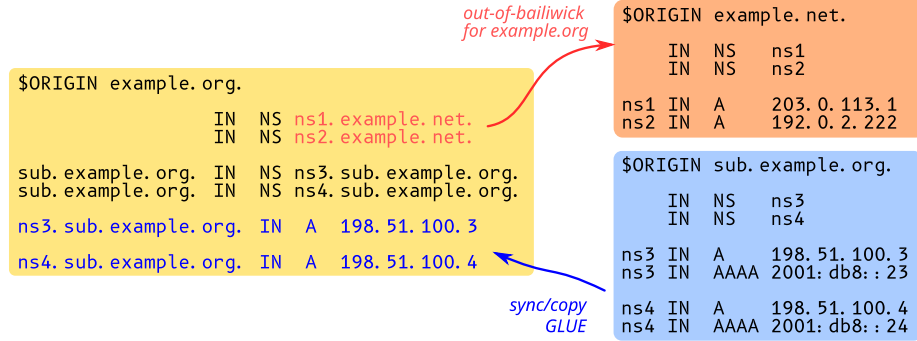


Fig. 1: Broken IPv6-delegation for `example.org` (missing AAAA resource records in `example.net` for NS) and `sub.example.org` (missing IPv6-GLUE in parent).

GLUE for IPv6, often is different. Hence, in this paper we use the term broken IPv6-delegation to avoid unnecessary ambiguity and distinguish the case of zones that are not IPv6 ready, e.g. show no intent to support IPv6 by not having any AAAA records, and zones that appear to intend supporting IPv6, Section 2.

In the first example, the external nameservers (“out-of-bailiwick”) of `example.org` do not have AAAA records and, thus, the resolution via IPv6 is impossible. In the second example, the zone `example.org` misses the AAAA glue records. These glue records make the A/AAAA records available for resolution if they have to be resolved from the zone being delegated, i.e., the names of the NS {ns3,ns4}.sub.example.org are in-bailiwick.

These examples highlight (a) that it needs cooperation between multiple parties for proper configuration, i.e., `sub.example.net` cannot be resolved via IPv6 even though it is correctly configured; (b) that dual-stack hides issues, i.e., both examples work for dual-stack enabled hosts where the AAAA records for ns3 and ns4 are resolvable. This demonstrates how working IPv4 resolution hides broken IPv6-delegation for dual-stack DNS recursors.

To be IPv6 *ready*, DNS resolution must work in IPv6-only scenarios. In this paper, we leverage passive DNS data—the Farsight SIE dataset [17]—to identify scenarios in which the DNS delegation chain breaks when only IPv6 is available. Our main contributions can be summarized as follows:

- We identify common broken IPv6-delegation scenarios and point out the importance of checking the full delegation chain.
- We show that big players have a major impact on the number of zones affected by broken IPv6-delegation. Today, 10 DNS providers are responsible for about 24.6% of IPv6-only-unresolvable domains we observe. Just by adding correct glue records, in Jan. 2017 one single provider fixed the IPv6-only name resolution of more than 45.6 M domains (15.5% of the domains in the dataset).
- Resilience mechanisms often hide misconfigurations. For example, broken IPv6-delegation is hidden by the combined efforts of DNS resilience and Happy

Eyeballs. Correctly configuring ones own DNS zone is not sufficient and dependencies are often non-obvious.

- Additionally, we conduct a thorough validation of our methodology. We assess the coverage of the Farsight SIE data in comparison to available ground-truth zonefile data, finding it to provide sufficient coverage for our analysis. Furthermore, we cross-validate our passive measurement results using active measurements, again finding our results to be robust.
- We implemented a DNS measurement tool instead of using, e.g., ZDNS [29], as we need IPv6 support which ZDNS does not (yet) support. The dataset from our active measurements and an implementation of our scanning methodology, including a single-domain version operators can use to evaluate IPv6 support for their own domains, are publicly available at:
<https://github.com/mutax/dns-v6-readiness>

2 Broken IPv6 Zone Delegation

In this section, we briefly recap DNS zone delegation, and sketch common DNS resolution failure scenarios.

2.1 Background: DNS Zone Delegation

The DNS is organized in a hierarchical structure where each node represents a zone that can be operated separately from its parent or child zones. For a zone to be resolvable, **NS** resource records have to be set in two places. First, the parent of the zone has to explicitly delegate the zone to authoritative nameservers via **NS** resource records. If an authoritative server has a domain name within the delegated zone itself or a child zone, i.e., if it is “in-bailiwick” [26], the parent zone must also contain **A** and **AAAA** resource records for this name, called **GLUE**, that are returned in the **ADDITIONAL** section of the DNS responses whenever the **NS** record is returned. This process breaks the circular dependency in the resolution chain. Furthermore, the zone itself must contain appropriate **NS** records as well as **A** and **AAAA** records if they are in-bailiwick. If the name in an **NS** record is not within the zone itself or a child zone, i.e., it is out-of-bailiwick, then the zone of the **NS**’ name must also resolve for the initial zone to be resolvable.

2.2 Reasons for Broken IPv6 Delegation

In this paper, we focus on a subset of DNS misconfigurations. In an IPv6-only scenario these misconfigurations can lead to effects similar to what is commonly referred to as lame delegation. To avoid ambiguity, we use the term broken IPv6-delegation referring to any set of misconfiguration specific to IPv6, that breaks the DNS delegation chain of a zone and prevents any of its records from resolving in an IPv6-only scenario. Other issues where a zone does not resolve due to, e.g., DNSSEC problems or unresponsive nameservers, i.e., the strict definition of “lame delegation” (see RFC8499 [26]) are out-of-scope. The issues we discuss

can also occur in IPv4 DNS resolution, but are usually quickly discovered given the currently still large number of sites with IPv4-only connection to the Internet, that will not be able to resolve the affected zones.

For a zone to be IPv6-resolvable —i.e., resolvable using IPv6-only— the zones of the authoritative nameservers have to be resolvable via IPv6 and at least one nameserver must be accessible via IPv6. This has to be the case *recursively*, i.e., not only for all parents of the zone itself but also for all parents of the authoritative nameservers in such a way that at least for one⁴ of the authoritative nameservers of a zone a delegation chain from the root zone exists, that is fully resolvable using IPv6. We identify the following misconfigurations which can cause broken IPv6-delegation in an IPv6-only setting:

- **No AAAA records for NS names:** If none of the NS records for a zone in their parent zone have associated AAAA records, resolution via IPv6 is not possible.
- **Missing GLUE:** If the name from an NS record for a zone is in-bailiwick, i.e., the name is within the zone or below [26], a parent zone must contain an IPv6 GLUE record, i.e., a parent must serve the corresponding AAAA record(s) as ADDITIONAL data when returning the NS record in the ANSWER section.
- **No AAAA record for in-bailiwick NS:** If an NS record of a zone points to a name that is in-bailiwick but the name lacks AAAA record(s) in its zone, IPv6-only resolution will fail even if the parent provides GLUE, when the recursive server validates the delegation path. One such example is Unbound [35] with the setting `harden-glue: yes`—the default.
- **Zone of out-of-bailiwick NSes not resolving:** If an NS record of a zone is out-of-bailiwick, the corresponding zone must be IPv6-resolvable as well. It is insufficient if the name pointed to by the NS record has an associated AAAA record.
- **Parent zone not IPv6-resolvable:** For a zone to be resolvable via IPv6 the parent zones up to the root zone must be IPv6-resolvable. Any non-IPv6-resolvable zone breaks the delegation chain for all its children.

The above misconfigurations are not mutually exclusive. For example, if the NS sets between parent and child differ, a common misconfiguration [42], the NS in the parent may not resolve due to missing GLUE (as they are in-bailiwick) *but also* the NS in the child may not resolve due to having no AAAA for their names, if they are out-of-bailiwick. In this paper we investigate the prevalence of these misconfigurations to evaluate the IPv6 readiness of the DNS ecosystem.

3 Datasets and Methodology

In this section, we present our choice of datasets as well as our active and passive measurement methodology for identifying DNS misconfigurations that break IPv6-only resolution.

⁴ RFC2182 [21] suggests to avoid such single points of failure

Table 1: List of data fields in the Farsight SIE dataset.

Field	Description	Example
count	# of times the tuple <code><rrname, rrtype, 12 bailiwick, rdata></code> has been seen.	
time_first	Unix timestamp of the first occurrence of the unique tuple during the data slice	1422251650
time_last	Unix timestamp of the last occurrence of the unique tuple during the data slice.	1422251650
rrname	Requested name in the DNS.	example.com
rrtype	Requested RRtype of the query.	NS
bailiwick	Zone authoritative for the reply.	com
rdata	List of all responses received in a single query.	["ns1.example.com", "ns2.example.com"]

3.1 DNS Dataset: Farsight SIE

For our evaluation we are looking for a dataset that enables us to (a) perform a longitudinal study, (b) detect IPv6 DNS misconfigurations, (c) analyze not just top level domains (TLDs) but *also* zones deeper in the tree, and (d) focus on zones that are used in-the-wild. As such we select the Farsight SIE dataset for our study.

The *Farsight Security Information Exchange* (SIE) dataset [17] is collected by Farsight Inc. via globally distributed sensors, co-located with recursive DNS resolvers. Each sensor collects and aggregates all DNS cache misses that the recursive DNS resolver encounters, i.e., the outgoing query and the received answer. By only recording cache-misses and providing aggregates, Farsight reduces the risk of exposing Personally Identifiable Information (PII). Cache-misses occur when a recursive DNS resolver does not have a DNS record for a specific domain name in its cache (or the record’s TTL has expired). The recursive resolver then has to ask the authoritative nameserver for the requested name, which is then recorded by Farsight SIE. Farsight does not share the exact number and location of its sensors for business confidentiality reasons. Farsight’s SIE dataset has been used in previous research [22, 27, 32] and its efficacy, coverage, and applicability for research has been demonstrated in the past [23]. We discuss ethical considerations of using this dataset in Section 3.5.

We use monthly aggregates from January 2015 to August 2022, containing unique tuples of: requested name, requested RRtype, bailiwick of the response, and returned data record, also for the additional sections, see Table 1. Thus, the Farsight dataset contains essential information for us, as it also records additional data as entries with the bailiwick of the parent. In addition, the Farsight dataset reaches deeper into the DNS hierarchy than, e.g., OpenINTEL [36], as it monitors DNS requests in the wild instead of resolving a set of names below zones sourced from TLD zone files.

Farsight Global Zone Coverage A common question when using a passive dataset like the one provided by Farsight is how well it actually covers zones on the Internet. In order to determine the coverage of the Farsight dataset, we evaluated the overlap of the second-level domains (SLDs) observed in the dataset with




Fig. 2: Zone coverage of Farsight data and number of zones used for the evaluation. We used available zone files to determine the share of covered second level domains by Farsight’s dataset. Please note the dip in the graph from February to August 2019, where our zone file collection was limited, i.e., we only collected few zones with high coverage (February - April and July, including .com), or no data at all (May and June).

ground-truth data, i.e., the names extracted from available zone files. Specifically, we are comparing to .com, .net, and other gTLD (generic TLD) zone files starting from mid of 2016. Additionally, from April 2017 onward, we also obtained CZDS (ICANN Centralized Zone Data Service) zone file data for all available TLDs. Moreover, we use publicly available zone file data from .se, .nu, and .ch for the coverage analysis. In total, this allows us to compare Farsight’s data to more than 1.1k zones as of August 2022.

Looking at coverage over time, we find a significant overlap between the Farsight dataset and the number of actually delegated zones based on zone files, see Figure 2. Coverage averages above 95 % from 2019 onwards, with especially since May 2021, our coverage reaches over 99 %. Furthermore, we find a reduced average coverage in the beginning of 2017. A closer investigation revealed that these relate to the introduction of various vanity gTLDs with an overall small size, i.e., below 100 delegated zones in the TLD. This implies that missing coverage for just a few zones would lead to a significant reduction in aggregate coverage. Nevertheless, our analysis shows that a significant share of zones is covered in the Farsight dataset. Hence, we the Farsight dataset—especially due to the historic perspective it provides—is ideal to investigate our research questions.

Despite this high coverage, we still face the drawback of the Farsight dataset relying on real-world usage. As such, a missing record in the passive dataset does not necessarily indicate non-existence. Hence, we independently corroborate all major findings with data from TLD zone files for a specific period to check for missing glue records in the zone file, see Section 5.4.

3.2 Domain Classification

There are many ways to cluster DNS domains into subgroups. For example, one may look only at the *Top Level Domains* as specified by ICANN [28], or use the *Public Suffix List (PSL)* provided by the Mozilla Foundation [34] to identify second level domains. The PSL is used by browser vendors to decide if a domain

is under private or public control, e.g., to prevent websites from setting a *super-cookie* for a *domain* such as *.co.uk*. Based on matching monthly samples of the ICANN TLDs and the PSLs we identify *TLDs* as well as *2nd Level Domains*, and *Zones Below 2nd Level*, i.e., all zones *below 2nd Level Domains*.

Another way of grouping DNS domains is to use the Alexa Top-1M list [3]. Using, again, matching monthly samples, we distinguish between the Top 1K, Top 1K–10K, Top 10K–100K, and Top 100K–1M domains. We note that there are limitations in the Alexa Top List [39,40], but compared to other toplists such as Tranco [31], the Alexa list is available throughout the measurement period.

3.3 Misconfiguration Identification

Here, we describe how we identify whether zones can be resolved only via IPv4, only via IPv6, via IPv4 and IPv6, or not at all from the dataset.

1. Per zone NS set identification: We first identify all zone delegations by extracting all entries with `rrtype = NS`. Next, for all names used in these delegations, we find all associated IPs by extracting all `A` and `AAAA` records. We do not consider `CNAMEs` since they are invalid for `NS` entries, see RFC2181 [20].

We then iterate over all zones, i.e., names that have `NS` records, to create a unique zone list. In this process, we record the `NS` records for each bailiwick sending responses for this zone observed in the dataset, and for each `NS` name all `AAAA` and `A` type responses, again grouped by bailiwick from which they were seen. This also captures cases where parent and child return *different* `NS` sets.

2. Per zone DNS resolution: We consider a zone to be resolvable via IPv4 or IPv6 if *at least one* of the `NS` listed for the zone can be resolved via IPv4 or IPv6 respectively. Hence, to check which zones can be resolved using which IP protocol version we simulate the DNS resolution, starting at the root, i.e. we assume the Root zone `.` to be resolvable by IPv4 and IPv6. We then iterate over the zone set with attached `NS` and `A/AAAA` data. For each zone, except the root zone, we initialize an empty state marking the zone as not resolving.

We then attempt to resolve each zone. For that, we first check if the zone’s parent has been seen⁵. If so we check for each `NS` of the zone we are trying to resolve as listed in the parent whether its name resolves via IPv4 and/or IPv6. This is the case if:

1. The `NS` is outside the zone we are trying to resolve, the `NS`’ zone has been recorded as resolving in the zone state file (via IPv4 and/or IPv6), and there are `A/AAAA` records with that zone’s bailiwick for the `NS`.
2. The `NS` is in the zone we are resolving and there is an `A/AAAA` glue record for the name with the bailiwick of the zone’s parent (only if an in-bailiwick `NS` is listed in the parent).

⁵ Due to an implementation oversight in our passive data evaluation, we only checked for the parent having been seen, not that it actually resolves via IPv4/IPv6. This may lead to us underestimating the number of non-resolving zones in the range of $\leq 0.1\%$ of all zones and $\leq 0.05\%$ of all second level domains.

Algorithm 1 Resolve Zones from Passive Data

```

1:  $zone\_res \leftarrow \{\}$ 
2:  $ns\_res \leftarrow \{\}$ 
3:  $prev\_res\_zones \leftarrow -1$ 
4:  $cur\_res\_zones \leftarrow 0$ 
5:
6: while  $!prev\_res\_zones == cur\_res\_zones$  do
7:    $prev\_res\_zones \leftarrow cur\_res\_zones$ 
8:    $cur\_res\_zones \leftarrow 0$ 
9:   for  $zone$  in  $input$  do
10:    if  $zone\_res[zone.parent][res]$  then
11:       $glue\_resolve \leftarrow false$ 
12:       $zone\_resolve \leftarrow false$ 
13:      for  $NS$  in  $glue$  do
14:        if  $NS$  in  $ns\_res$  ||  $(NS$  in  $zone$  &&  $zone.parent$  has  $NS.ip$ ) ||
15:         $(zone\_res[ns\_zone][res]$  &&  $ns\_zone$  has  $NS.ip)$  then
16:           $ns\_res[NS] \leftarrow true$ 
17:           $glue\_resolve \leftarrow true$ 
18:        for  $NS$  in  $zone$  do
19:          if  $NS$  in  $ns\_res$  ||  $(NS$  in  $zone$  &&  $zone$  has  $NS.ip)$  ||
20:           $(zone\_res[ns\_zone][res]$  &&  $ns\_zone$  has  $NS.ip)$  then
21:             $ns\_res[NS] \leftarrow true$ 
22:             $zone\_resolve \leftarrow true$ 
23:           $zone\_res[zone][glue\_res] \leftarrow glue\_resolve$ 
24:           $zone\_res[zone][zone\_res] \leftarrow zone\_resolve$ 
25:          if  $glue\_resolve$  &&  $zone\_resolve$  then
26:             $zone\_res[zone][res] \leftarrow true$ 
27:           $cur\_res\_zones \leftarrow cur\_res\_zones + 1$ 

```

To ensure full resolution, we also have to check that the NS listed in the child resolve. For NS with names under the zone this is the case if the NS listed for this zone in the parent can be reached via IPv4/IPv6, see above, and they have A/AAAA records with the bailiwick of the zone itself. For out-of-bailiwick NS, this is again the case if their own zone resolves and they have A/AAAA records.

A single iteration of this process is not sufficient, as zones often rely on out-of-bailiwick NS. Hence, we continue iterating through the list of zones until the number of unresolved zones no longer decreases. For a simplified pseudo-code description, see Algorithm 1.

3.4 Active Measurement Methodology

To validate our passive measurement results, we implemented a resolver in python. While, technically, Izhikevich et al. presented ZDNS, a tool for this purpose, ZDNS does not provide sufficient support for IPv6 resolution for our use-case. Our measurement methodology follows essentially the same algorithm as our passive resolution. For each zone, we start at the root, and iterate through

the DNS tree. From there, we query all authoritative nameservers recorded in the parent on each layer of the DNS hierarchy using IPv4 and IPv6 where possible for the NS of that zonelayer. Furthermore, we try to obtain any possibly available GLUE (A and AAAA) for in-bailiwick NS. For out-of-bailiwick NS, we try to resolve the NS, again starting from the root. If there is an inconsistency between parent and child, i.e., if we discover additional NS when querying the NS listed in the parent, we also perform all queries for this layer against these, noting that they were only present in the child.

To limit the amount of queries sent to each server, our implementation follows the underlying principles of QNAME minimization as described in RFC7816 [5]. By using the NS resource record type to query the parent zones we can directly infer zonecuts and store GLUE records from the additional section, if present. Note that RFC8020 [6] is still not implemented by all nameservers, thus we cannot rely on NXDOMAIN answers to infer that no further zones exist below the queried zone. Our measurement tool will retry queries using TCP on truncation and disable EDNS when it receives a FORMERR from the upstream server.

To further limit the number of queries sent, all responses, including error responses or timeouts, are cached. We limit the number of retries (4) as well as the rate (20 second wait time) at which they are sent. To further enrich the actively collected dataset, we query all authoritative nameservers of a zone for the NS, TXT, SOA and MX records of the given zone as well as the version of the used server software using the *version.bind* in the CHAOS class. Queries and replies are recorded tied to the NS that provided them.

We ran these measurements between October 10th to 14th and 22nd to 24th 2022 against the Alexa Top1M from August 15th 2022 containing 476,242 zones, collecting responses to a total of 32M queries sent via IPv4 and 24M queries sent via IPv6. Our active measurement dataset (101GB of json data), and a tool implementing our measurement toolchain are publicly available at: <https://github.com/mutax/dns-v6-readiness>

3.5 Ethical Considerations

The *Farsight Security Information Exchange* (SIE) dataset [17] used in this work is collected by Farsight Inc. at globally distributed vantage points, co-located to recursive DNS resolvers. These sensors collect and aggregate DNS cache misses they encounter, i.e., outgoing queries of the recursors and the received answers. Only collecting cache misses is a conscious choice by Farsight to ensure PII is protected. The dataset also does not contain which sensors collected a specific entry. We specifically use a per-month aggregated version of the dataset, see Section 3.1. For details on the fields in the dataset, see Table 1. Data has been handled according to best practices in network measurement data handling as outlined by Allman and Paxson [2].

Before running the active measurements for validation purposes (cf. Section 3.4), we consult the Menlo report [30] as well as best measurement practices [19]. We limit our probing rate, send only well-formed DNS requests, and make use of dedicated servers which have informative rDNS names set. Addition-

ally, we run a webserver providing additional information and contact details on the IP as well as on the rDNS name. We also focused our measurements on the Alexa Top 1M, i.e., sites for which the impact of additional requests at the scale of our measurements is not significant, while also limiting repeated requests using caching. During our active measurements, we did not receive any complaints. In summary, we conclude that this work does not raise any ethical issues.

4 Results

Here, we first provide an aggregate overview of the Farsight dataset. Subsequently, we present the results of our analysis of broken IPv6-delegation based on passive measurement data. Finally, we validate our passive measurement results against active measurements run from 10th to 24th of October 2022.

4.1 Dataset Overview

Our passive dataset spans 7 years starting on January 1st, 2015 and ending on August 31st, 2022. During this period, the number of unique zones increased from 126 M to 368 M. Similarly, the number of PSL 2nd level domains increased from 116 M to 326 M. For a visualization see the gray line in Figure 3 (right y-axis). To highlight our findings, we present results for selected subsets of domains only. The full results for all domain subsets are in shown in Appendix A.

4.2 IPv6 Resolution in DNS Over Time

In Figure 3 we show how the fraction of zones that is resolvable via IPv4-only, IPv6-only, both protocols, or fails to resolve, changes across time. We also show how the total number of zones changes (gray line). The figure shows data for all zones, the ICANN TLDs, PSL 2nd domains, zones deeper in the tree, Alexa Top-1K and Alexa Top-1M.

Overall, see Figure 3a, we find that 11.6% of all zones are IPv6-resolvable in January 2015. This is significantly higher than the sub 1% reported by Czyz et al. [13] in 2014. However, they only accounted for glue records, which does not consider zones with out of bailiwick NS. Over time IPv6 adoption steadily increases, with 55.5% of zones resolving via IPv6 in August 2022. A notable increase of IPv6 resolvable zones by 17.4% occurs in January 2017. Further investigation we find, that this increase relates to two major DNS providers—a PaaS provider and a webhoster—adding AAAA glue for their NS.

For ICANN TLDs, see Figure 3b, we find that the majority of zones is IPv6-resolvable. A notable increase—2.94%—occurs in January 2018. We attribute the increase in resolvability to a popular search engine operator, holding several vanity TLDs, enabling IPv6 resolution by adding IPv6 GLUE. By the end of our observation period nearly all TLDs are IPv6-resolvable. The remaining not IPv6-resolvable zones are several vanity TLDs as well as smaller ccTLDs.

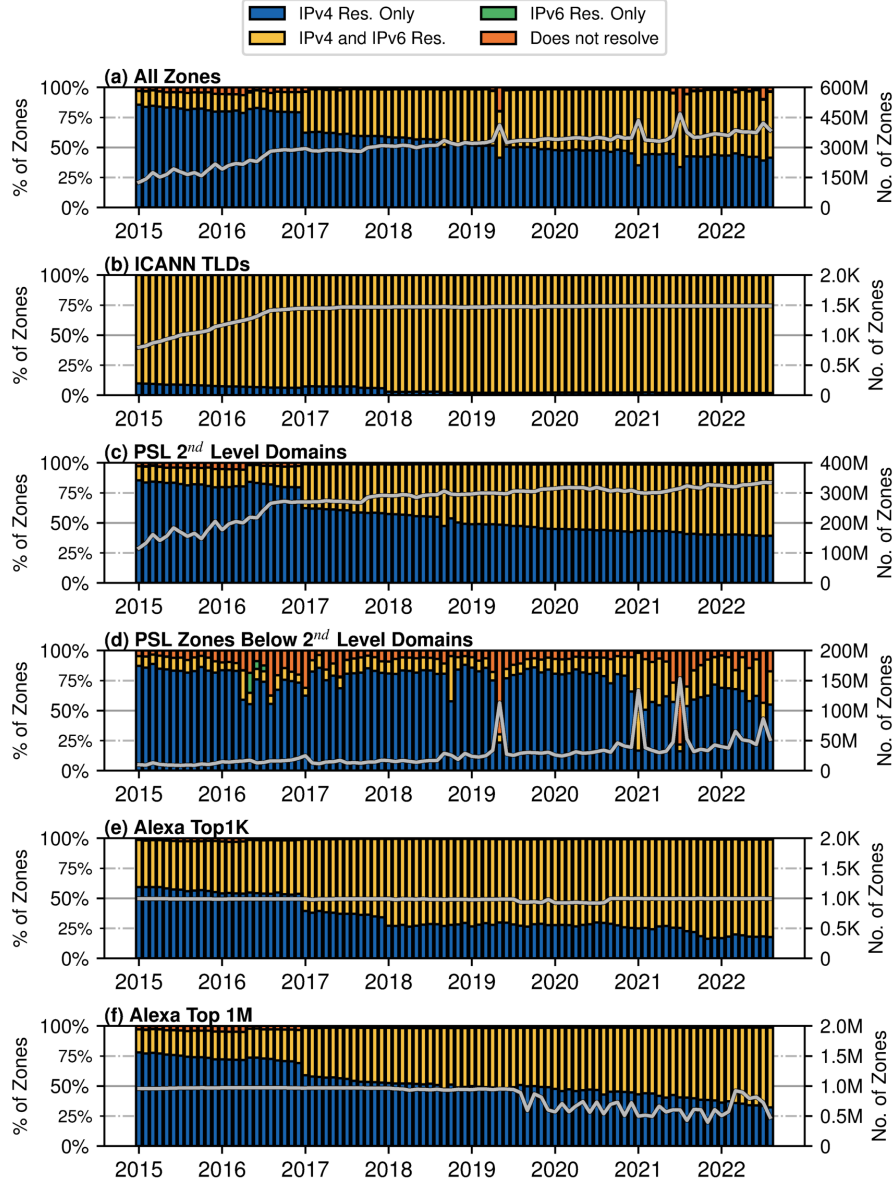


Fig. 3: Per month: # of zones (gray line—right y-axis) and IPv4/IPv6 resolvability in % (left y-axis).

While PSL 2^{nd} level domains, see Figure 3c, mirror the general trend of all zones, we find that zones deeper in the tree (Figure 3d) are generally less likely to be IPv6-resolvable. Still, we observe an upward trend. We attribute this to the fact that the process of entering such domains into TLDs for 2^{nd} level domains still receives oversight by NICs, e.g., regarding the RFC compliant use of at least two NS in different networks [21], while zones below 2^{nd} level domains can be freely delegated by their domain owners. Also, for sub-domains, we observe three distinct spikes in Figure 3d which correspond to the spikes seen for all domains, recall Figure 3a. These spikes occur when a single subtree of the DNS spawns millions of zones. These are artifacts due to specific configurations and highlight that lower layer zones may not be representative for the overall state of DNS.

Finally, comparing PSL 2^{nd} level domains, see Figure 3d, to the Alexa Top-1K domains, see Figure 3e, we find that IPv6 adoption is significantly higher among popular domains, starting from 39.2% in 2015 and rising to 82.6% in 2021. There are two notable steps in this otherwise gradual increase, namely January 2017 and January 2018. These are due to a major webhoster and a major PaaS provider enabling IPv6 resolution (2017), and a major search engine provider common in the Alexa-Top-1K enabling IPv6 resolution (2018).

Comparison with Active Measurements: Evaluating zone resolvability from our active measurements, see Section 3.4, we find that 314,994 zones (66.14%) support dual stack DNS resolution, while 159,166 zones (33.42%) are only resolvable via IPv4.

A further 2066 zones (0.43%) could not be resolved during our active measurements, and 16 zones ($\leq 0.01\%$) were only resolvable via IPv6. In comparison to that, our passive measurements—see also Figure 3f—map closely: We find 66.36% (+0.22% difference) of zones in the Alexa Top 1M resolving via both, IPv4 and IPv6, and 32.21% (-1.21% difference) of zones only resolving via IPv4. Similarly, 1.01% (+0.58%) of zones do not resolve at all, and 0.43% (+0.43% difference) of zones only resolve via IPv6 according to our passive data. Hence, overall, we find our passive approach being closely aligned with the results of our active measurements for the latest available samples. The, in comparison, higher values for non-resolving and IPv6 only resolving zones are most likely rooted in the visibility limitations of the dataset, see Section 5.4. Nevertheless, based on the low deviation between two independent approaches at determining IPv6 resolvability of zones we have confidence in the results of our passive measurements.

4.3 IPv6 Resolution Failure Types

Next, we take a closer look at zones that show some indication of IPv6 deployment, yet, are not IPv6-resolvable. These are zones where an NS has an AAAA record or an AAAA GLUE. To find them we consider NS entries within the zone as well as NSes for the zone in its parent. In Figure 4 we show how their absolute numbers evolve over time (gray line) as well as the failures reasons (in percentages).

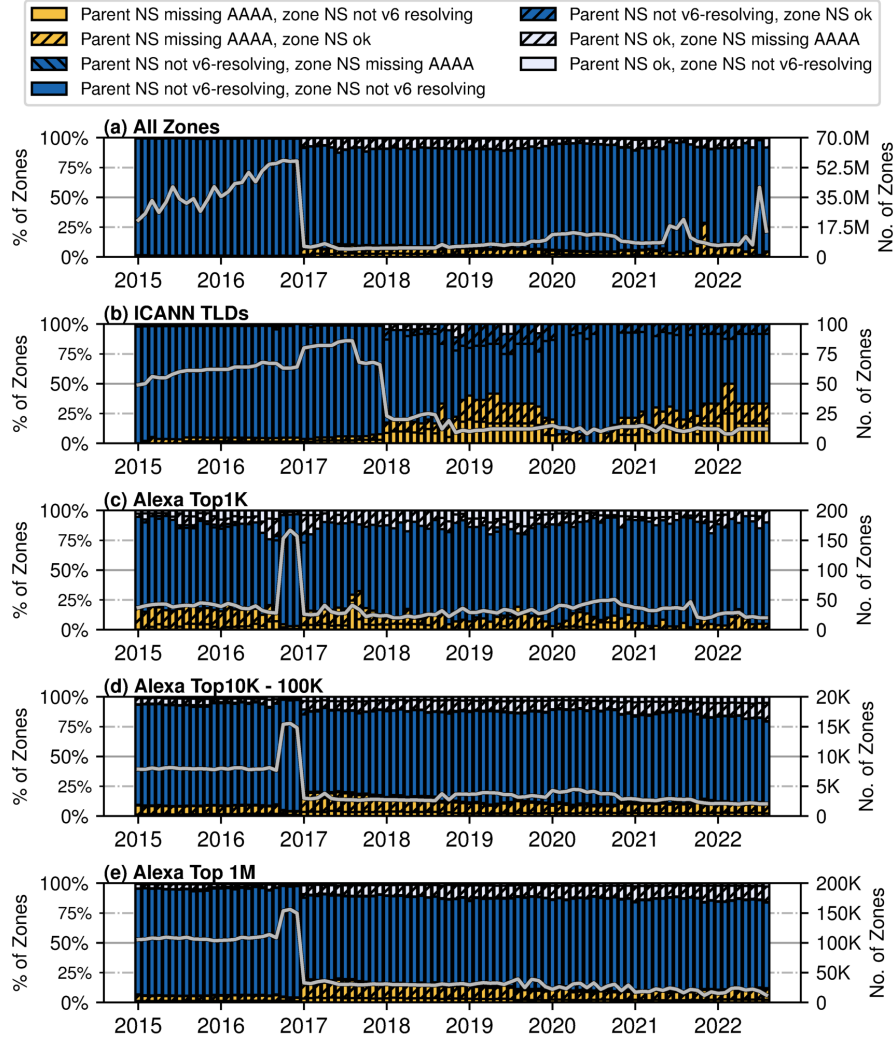


Fig. 4: Per month: # of zones not IPv6-resolvable with AAAA or GLUE for NS (gray line-right y-axis) and causes for IPv6 resolution failure in % (left y-axis).

We find that for all four subsets of zones shown—all zones, ICANN TLDs, Alexa Top-1K, Alexa Top-10K–100K—the most common failure case is missing resolution of NS in the parent. This occurs mostly when the NS is out-of-bailiwick and *does* have AAAA records, but the NS’s zone itself is not IPv6-resolvable. Furthermore, there is a substantial number of zones per category—especially in the Alexa Top-1K—where the NS in the parent lacks AAAA while the NS listed in the zone has AAAA records, commonly due to missing GLUE. We also observe the inverse scenario, i.e., GLUE is present but no AAAA record exist for the NS within the zone itself. Both cases can also occur if NS sets differ between the parent and its child [42].

We see a major change around January 2017, i.e., a sharp increase in zones that are IPv6-resolvable, which is also visible in Figure 3: For all zones as well as for the Alexa Top 10K–100K, we observe that several million zones not resolving via IPv6 since the start of the dataset but having NSes with AAAA records, now are IPv6-resolvable. The reason is that a major provider added missing glue records. Interestingly, we do not see this in the Alexa Top 1K.

In the Alexa Top 1K, and to a lesser degree in the Alexa Top 10K–100K, we observe a spike of zones that list AAAA records for their NS but are not IPv6-resolvable in Oct. 2016. This is the PaaS provider mentioned before, first rolling out AAAA records for their NS, and then three months later also adding IPv6 GLUE. Operationally, this approach makes sense, as they can first test the impact of handling IPv6 DNS queries in general. Moreover, reverting changes in their own zones is easier than reverting changes in the TLD zones—here the GLUE entries. Again, the major webhoster is less common among the *very* popular domains, which is why its effect can be seen in Figures 4a and 4d, but not in Figure 4b. Also, this operator had AAAA records in place since the beginning of our dataset, as seen by the plateau in Figure 4d. These observations have been cross-confirmed by inspecting copies of zonefiles for the corresponding TLDs and time-periods.

4.4 Centralization and IPv6 Readiness

Finally, we focus on the nameservers hosting most non IPv6-resolvable zones. We first identify the top NS sets in terms of the number of hosted zones, aggregating NS names to their PSL 2nd level domain and known operators’ NS under a multiple well-patterned zones. Then, we compute a CDF over the number of zones per NS set for each time bin. Figure 5 shows how this CDF changed across time and highlights the impact of centralization within the DNS providers. Over 97.5% of the non-IPv6-resolvable zones are hosted by the Top 10% of NS sets.

Again, we see the impact of a change by a major webhoster in January 2017—it is the top NS set among all zones (Figure 5a). Similarly, the PaaS provider is pronounced among the Alexa Top-1K, i.e., part of the Top 10 of NS sets (Figure 5c) and the top NS set for the Alexa Top-10K–100K (Figure 5d). Finally, the major search engine operator’s impact can especially be seen among TLDs (Figure 5b) and the Alexa Top-1K (Figure 5c), where—in both cases—this operator is the top NS set for non IPv6-resolvable zones.

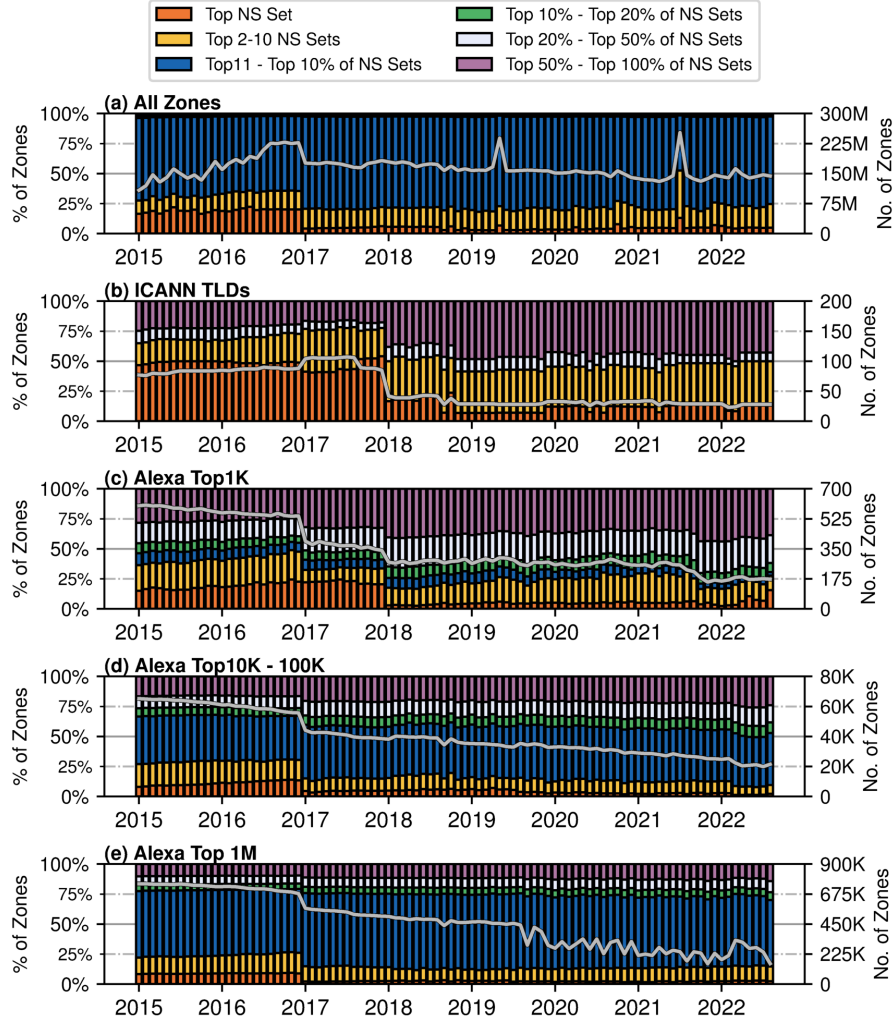


Fig. 5: Per month: # of zones not IPv6-resolvable (gray line-right y-axis) and distribution of zones over NS sets in % (left y-axis).

4.5 Resolvability and Responsiveness of NS in Active Measurements

During our active measurements, we also had the opportunity to validate whether NS records listed in zones did actually reply to DNS requests or not. During our evaluation of the Alexa Top 1M, we discovered a total of 176,207 NS records, of which 212 had A or AAAA records associated that were invalid, as for example :: as a AAAA record. Of the remaining 175,995 records, 116,504 needed glue, i.e., they were in-bailiwick NS for their own. Among these, 19,310 NS were dual-stack, while 94,192 only had A records associated with them, and a further 108 NS only had associated AAAA records. Furthermore, 85,213 (90.47%) of A-only NS needing glue had correct glue set. For dual-stack configured NS, 14,072 (72.87%) have complete (A and AAAA) glue. A further 3,932 (20.36%) NS only has A glue records, while 24 (0.12%) NS only have AAAA glue, despite generally having a dual-stack DNS configuration. Finally, of the 108 NS records only having AAAA records associated, 70 (64.81%) NS have correctly set AAAA glue.

Moving on to the reachability of these NS, we find that of the total number of NS that have an A record (169,547) *and* are reachable is at 164,255, or 96.88% actually responds to queries. For IPv6, these values are slightly worse, with 30,193 of 32,285 NS (93.52%) responding to queries via IPv6. This highlights a potential accuracy gap of 3-6% for research work estimating DNS resolvability from passive data. Notably, this gap is larger for IPv6.

5 Discussion

In this section, we first state our key-findings, and then discuss their implications.

5.1 The Impact of Centralization

Centralization is one of the big changes in the Internet over the last decade. This trend ranges from topology flattening [4,7] to the majority of content being served by hypergiants [8] and—as we show—also applies to the DNS. An increasing number of zones are operated by a decreasing number of organizations. As such, an outage at one big DNS provider [44]—or missing support for IPv6—can disrupt name resolution for a very large part of the Internet as we highlight in Section 4. In fact, out-of-bailiwick NS not being resolvable via IPv6 is the most common misconfiguration in our study, often triggered by missing GLUE in a single zone. Given that *ten* operators could enable IPv6 DNS resolution for 24.6% of not yet IPv6 resolving zones, we claim that large DNS providers have a huge responsibility for making the Internet IPv6 ready.

5.2 IPv6 DNS Resolution and the Web

In general, as we travel down the delegation chain we find more misconfigurations and a smaller fraction of IPv6-resolvable zones. Given that common web assets—JavaScript, Style Sheets, or images—are often served from FQDNs further down the DNS hierarchy, we conjecture that this may have a another huge, yet still hidden, impact on IPv6 readiness for web. We encourage operators to be mindful of this issue, and study its effect in future work.

5.3 Implications for Future Research

Our findings demonstrate that it is not sufficient to test for the presence of **AAAA** records to assess the IPv6 readiness of a DNS zone. Instead, measurements have to assess whether the zones are IPv6-resolvable. The same applies to email setups and websites.

Furthermore, given the centralization we observe in the DNS, network measurements of IPv6 adoption should consider and quantify the impact of individual operators. More specifically, researchers should distinguish between effects caused by a small number of giants vs. the behavior of the Internet at large. Artifacts that can occur temporarily should be recognized and then excluded.

5.4 Limitations

Since our dataset relies on DNS cache misses, we are missing domains that are not requested or not captured by the Farsight monitors in a given month. Moreover, our use of monthly aggregates may occlude short-term misconfigurations. To address this, we support major findings on misconfigurations with additional ground-truth data from authoritative TLD zone files.

Similarly, we use the Alexa List with its known limitations [39, 40]. Thus, we cluster the Alexa list into different rank tiers, which reduces fluctuations in the higher tiers. Furthermore, we only assess zones' configuration states, and not actual resolution, i.e., "lame delegation" for other reasons is out of scope.

Furthermore, we cannot make statements on whether the zones we measure *actually* resolve, e.g., if there is an authoritative DNS server listening on a configured IP address returning correct results. Still, we have certainty that zones we measure as resolvable are at least sufficiently configured for resolution. Similarly, we can not assess the impact of observed DNS issues on other protocols, e.g., HTTPs. To further address this limitation of our passive data source, we conducted active measurements, which validated the observations from our passive results and added further insights on the actual reachability of authoritative DNS servers for zones.

Naturally, our active measurements also have several limitations that have to be recorded. First, we conducted our measurements from a single vantage point. Given load balancing in CDNs via DNS [43], this may have led to a vantage point specific perspective. Nevertheless, we argue that misconfigurations [14] are likely to be consistent across an operator, i.e., the returned **A** or **AAAA** records may change, but not the issue of, e.g., missing **GLUE**. Furthermore, DNS infrastructure tends to be less dynamic than **A** and **AAAA** records.

Second, our measurements were only limited to the Alexa Top 1M and associated domains. We consciously made this choice instead of, e.g., running active measurements on *all* zones in the Farsight dataset to reduce our impact on the Internet ecosystem.

In summary, our study provides an important first perspective on IPv6 only resolvability. We suggest to complement our study with active measurements of IPv6 only DNS resolution and the impact of broken IPv6-delegation on the IPv6 readiness of the web due to asset dependencies as future work.

6 Related work

Our related work broadly clusters into two segments: *i)* Studies on IPv6 adoption and readiness, and *ii)* Studies about DNS and DNS misconfigurations.

6.1 IPv6 Adoption and Readiness

With the exhaustion of the IPv4 address space [38], IPv6 adoption has been a frequent topic of study. In 2014, Czyz et al. [13] conducted a primer study on IPv6 adoption, taking a multi-perspective approach that also covered DNS. Our measurements shed light on the time after their measurements which concluded in 2014. Furthermore, they estimate IPv6 adoption in DNS by only surveying AAAA glue records in `net.` and `com.`, while we consider the full resolution path. Work by Foremski et al. [24] and Plonka & Berger [37] investigate IPv6 adoption at the edge, which is orthogonal to our work. In recent years, various researchers took country and domain specific perspectives on IPv6 adoption, e.g., [12, 25, 33].

6.2 DNS and DNS Misconfiguration Studies

Since DNS is a core component of the Internet, it has been studied regularly over the past decades, including studies regarding the adoption of new protocol features, e.g., [9–11, 15, 16, 43]. Such studies use various active datasets, e.g., OpenINTEL [36], as well as passive datasets, e.g., the Farsight SIE dataset which we rely on, to, e.g., study operational aspects of the DNS [23]. More specifically focusing on DNS (mis)configuration, Sommesse et al. [42] study inconsistencies in parent and child NS sets and Akiwate et al. [1] work on lame delegation. However, contrary to our work, the latter two either do not consider the IP part of DNS delegation (Sommese et al.), or explicitly focus on IPv4 (Akiwate et al.). More recently, Izhikevich et al. presented ZDNS, a tool for large-scale studies of the DNS ecosystem in the Internet [29]. Unfortunately, ZDNS is tailored towards IPv4 and does not support querying authoritative nameservers over IPv6. Therefore, we cannot make use of ZDNS in our study. Instead we perform active DNS measurements with our own implementation of a DNS resolution methodology, which implements IPv6 resolution.

6.3 Summary

We expand on earlier contributions regarding IPv6 adoption. We provide a more recent perspective on the IPv6 DNS ecosystem and take a more complete approach to assess the IPv6 readiness in an IPv6-only scenario. This focus on IPv6 is also our novelty in context to earlier work on DNS measurements and DNS misconfigurations, which did not focus on how IPv6 affects DNS resolvability. Additionally, our active measurements for validating our passive measurement results also highlight that the presence of AAAA records does not necessarily imply IPv6 resolvability. Instead, to measure IPv6 resolvability, the resolution state of provided IPv6 resources has to be validated.

7 Conclusion

In this paper, we present a passive DNS measurement study on root causes for broken IPv6-delegation in an IPv6 only setting. While over time we see an increasing number of zones resolvable via IPv4 and IPv6, in August 2022 still 44.6% are not resolvable via IPv6. We identify not resolvable **NS** records of the zone or its parent as the most common failure scenario. Our recommendations to operators include to explicitly monitor IPv6 across the entire delegation chain.

Additionally, we conducted a dedicated validation of our results using active measurements. This validation broadly confirmed our results from the passive measurements and further highlighted the importance of not only relying on the presence of specific records, as nameservers for which IPv6 addresses are listed in the DNS may not actually be responsive.

We plan to provide an open-source implementation of our measurement methodology along with the paper. Furthermore, we will provide a reduced implementation of our measurement toolchain which will enable operators to explicitly check a given zone or FQDN for IPv6-resolvable. Similarly, we will provide the results of our active measurements as open data.

For future work we suggest to systematically expand our active measurement campaign to assess resolvability, e.g., for websites including all web assets. Using active measurements, one can explicitly resolve a hostname and run active checks on the delegation chain, validating the responses of all authoritative nameservers and find inconsistencies not only between a zone and its parent but also within the NS set. We conjecture that—especially given the widespread use of subdomains for web assets—the reduced IPv6 resolvability we observe may have a significant impact on the IPv6-readiness of the web, i.e., a website using assets on domains that do not resolve via IPv6 is not IPv6 ready.

Acknowledgments

We thank Farsight Security, Inc. (now DomainTools) for providing access to the Farsight Security Information Exchange’s passive DNS data feed. Without this data, the project would not have been possible. The authors express their gratitude to the anonymous reviewers for their thoughtful and encouraging input during the reviewing process. This work was partially funded by the German Federal Ministry of Education and Research under the project PRIME-net, grant 16KIS1370, and 6G-RIC, grant 16KISK027. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Farsight Security, Inc., DomainTools, the German Federal Ministry of Education and Research or the authors’ host institutions and further affiliations.

References

1. Akiwate, G., Jonker, M., Sommesse, R., Foster, I., Voelker, G.M., Savage, S., Claffy, K.: Unresolved Issues: Prevalence, Persistence, and Perils of Lazy Delegations. In: Proc. of the Internet Measurement Conference (IMC). p. 281–294. ACM (2020). <https://doi.org/10.1145/3419394.3423623>
2. Allman, M., Paxson, V.: Issues and Etiquette Concerning Use of Shared Measurement Data. In: Proc. of the Internet Measurement Conference (IMC). pp. 135–140. ACM (2007). <https://doi.org/10.1145/1298306.1298327>
3. Amazon.com, Inc: Alexa Top Sites. <https://www.alexa.com/>
4. Arnold, T., He, J., Jiang, W., Calder, M., Cunha, I., Giotsas, V., Katz-Bassett, E.: Cloud provider connectivity in the flat Internet. In: Proc. of the Internet Measurement Conference (IMC). pp. 230–246. ACM (2020). <https://doi.org/10.1145/3419394.3423613>
5. Bortzmeyer, S.: DNS Query Name Minimisation to Improve Privacy. RFC 7816 (Experimental) (Mar 2016), <https://www.rfc-editor.org/rfc/rfc7816.txt>, obsoleted by RFC 9156
6. Bortzmeyer, S., Huque, S.: NXDOMAIN: There Really Is Nothing Underneath. RFC 8020 (Proposed Standard) (Nov 2016), <https://www.rfc-editor.org/rfc/rfc8020.txt>
7. Böttger, T., Antichi, G., Fernandes, E.L., di Lallo, R., Bruyere, M., Uhlig, S., Castro, I.: Shaping the Internet: 10 Years of IXP Growth. arXiv (2019). <https://doi.org/10.48550/ARXIV.1810.10963>, <https://arxiv.org/abs/1810.10963>
8. Böttger, T., Cuadrado, F., Tyson, G., Castro, I., Uhlig, S.: A Hypergiant’s View of the Internet. ACM Computer Communication Review (CCR) **47**(1) (2017)
9. Calder, M., Fan, X., Hu, Z., Katz-Bassett, E., Heidemann, J., Govindan, R.: Mapping the Expansion of Google’s Serving Infrastructure. In: Proc. of the Internet Measurement Conference (IMC). p. 313–326. ACM (2013). <https://doi.org/10.1145/2504730.2504754>
10. Chhabra, R., Murley, P., Kumar, D., Bailey, M., Wang, G.: Measuring DNS-over-HTTPS Performance around the World. In: Proc. of the Internet Measurement Conference (IMC). p. 351–365. ACM (2021). <https://doi.org/10.1145/3487552.3487849>
11. Chung, T., van Rijswijk-Deij, R., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Wilson, C.: Understanding the Role of Registrars in DNSSEC Deployment. In: Proc. of the Internet Measurement Conference (IMC). p. 369–383. ACM (2017). <https://doi.org/10.1145/3131365.3131373>
12. Colitti, L., Gunderson, S.H., Kline, E., Refice, T.: Evaluating IPv6 Adoption in the Internet. In: Proc. of the Passive and Active Measurement Conference (PAM). pp. 141–150. Lecture Notes in Computer Science, Springer (2010)
13. Czyz, J., Allman, M., Zhang, J., Iekel-Johnson, S., Osterweil, E., Bailey, M.: Measuring IPv6 Adoption. In: Proc. of the 2014 ACM SIGCOMM Conference (SIGCOMM). p. 87–98. ACM (2014). <https://doi.org/10.1145/2619239.2626295>
14. Dietrich, C., Krombholz, K., Borgolte, K., Fiebig, T.: Investigating system operators’ perspective on security misconfigurations. In: Proc. of the 25th ACM SIGSAC Conference on Computer and Communications Security (CCS). pp. 1272–1289. ACM (2018)
15. Doan, T.V., Fries, J., Bajpai, V.: Evaluating Public DNS Services in the Wake of Increasing Centralization of DNS. In: IFIP Networking Conference (2021). <https://doi.org/10.23919/IFIPNetworking52078.2021.9472831>

16. Doan, T.V., Tsareva, I., Bajpai, V.: Measuring DNS over TLS from the Edge: Adoption, Reliability, and Response Times. In: Hohlfeld, O., Lutu, A., Levin, D. (eds.) *Proc. of the Passive and Active Measurement Conference (PAM)*. pp. 192–209. *Lecture Notes in Computer Science*, Springer (2021)
17. DomainTools, formerly Farsight Security: Farsight Security Information Exchange (SIE). <https://www.farsightsecurity.com/solutions/security-information-exchange/> (2022)
18. Durand, A., Ihren, J.: DNS IPv6 Transport Operational Guidelines. RFC 3901 (Best Current Practice) (Sep 2004), <https://www.rfc-editor.org/rfc/rfc3901.txt>
19. Durumeric, Z., Wustrow, E., Halderman, J.A.: ZMap: Fast Internet-wide Scanning and Its Security Applications. In: *Proc. of the 31th USENIX Security Symposium (USENIX Security)*. pp. 605–620. USENIX Association (2022)
20. Elz, R., Bush, R.: Clarifications to the DNS Specification. RFC 2181 (Proposed Standard) (Jul 1997), <https://www.rfc-editor.org/rfc/rfc2181.txt>, updated by RFCs 4035, 2535, 4343, 4033, 4034, 5452, 8767
21. Elz, R., Bush, R., Bradner, S., Patton, M.: Selection and Operation of Secondary DNS Servers. RFC 2182 (Best Current Practice) (Jul 1997), <https://www.rfc-editor.org/rfc/rfc2182.txt>
22. Fiebig, T., Borgolte, K., Hao, S., Kruegel, C., Vigna, G.: Something from nothing (There): collecting global IPv6 datasets from DNS. In: *Proc. of the Passive and Active Measurement Conference (PAM)*. *Lecture Notes in Computer Science*, vol. 10176, pp. 30–43. Springer (2017)
23. Foremski, P., Gasser, O., Moura, G.C.: DNS observatory: The big picture of the DNS. In: *Proc. of the Internet Measurement Conference (IMC)*. pp. 87–100. ACM (2019)
24. Foremski, P., Plonka, D., Berger, A.: Entropy/IP: Uncovering Structure in IPv6 Addresses. In: *Proc. of the Internet Measurement Conference (IMC)*. p. 167–181. ACM (2016). <https://doi.org/10.1145/2987443.2987445>
25. Han, C., Li, Z., Xie, G., Uhlig, S., Wu, Y., Li, L., Ge, J., Liu, Y.: Insights into the issue in IPv6 adoption: A view from the Chinese IPv6 Application mix. *Concurrency and Computation: Practice and Experience* **28**(3), 616–630 (2016). <https://doi.org/https://doi.org/10.1002/cpe.3327>
26. Hoffman, P., Sullivan, A., Fujiwara, K.: DNS Terminology. RFC 8499 (Best Current Practice) (Jan 2019), <https://www.rfc-editor.org/rfc/rfc8499.txt>
27. Houser, R., Hao, S., Li, Z., Liu, D., Cotton, C., Wang, H.: A Comprehensive Measurement-based Investigation of DNS Hijacking. In: *Proc. of the 40th International Symp. on Reliable Distributed Systems (SRDS)*. pp. 210–221. IEEE (2021)
28. ICANN: List of Top-Level Domains. <https://www.icann.org/resources/pages/tlds-2012-02-25-en>
29. Izhikevich, L., Akiwate, G., Berger, B., Drakontaidis, S., Ascherman, A., Pearce, P., Adrian, D., Durumeric, Z.: ZDNS: A Fast DNS Toolkit for Internet Measurement. In: *Proc. of the Internet Measurement Conference (IMC)*. ACM (2022)
30. Kenneally, E., Dittrich, D.: The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research. Available at SSRN 2445102 (2012)
31. Le Pochat, V., Van Goethem, T., Tajalizadehkhoob, S., Joosen, W.: TRANCO: A Research-Oriented Top Sites Ranking Hardened Against Manipulation. In: *Proc. of the 26th Network and Distributed System Security Symposium (NDSS)*. Internet Society (ISOC) (2019)

32. Liu, B., Lu, C., Li, Z., Liu, Y., Duan, H.X., Hao, S., Zhang, Z.: A Reexamination of Internationalized Domain Names: The Good, the Bad and the Ugly. In: Proc. of the 48th IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). pp. 654–665. IEEE (2018)
33. Livadariu, I., Elmokashfi, A., Dhamdhare, A.: Measuring IPv6 Adoption in Africa. In: e-Infrastructure and e-Services for Developing Countries. pp. 345–351. Springer International Publishing, Cham (2018)
34. Mozilla Foundation: Public Suffix List. <https://publicsuffix.org/>
35. NLnet Labs: Unbound Nameserver Documentation. <https://unbound.docs.nlnetlabs.nl/en/latest/reference/history/requirements.html>
36. OpenINTEL project: The OpenINTEL measurement platform. <https://openintel.nl/>
37. Plonka, D., Berger, A.: Temporal and Spatial Classification of Active IPv6 Addresses. In: Proc. of the Internet Measurement Conference (IMC). p. 509–522. ACM (2015). <https://doi.org/10.1145/2815675.2815678>
38. Richter, P., Allman, M., Bush, R., Paxson, V.: A Primer on IPv4 Scarcity. ACM Computer Communication Review (CCR) **45**(2), 21–31 (2015). <https://doi.org/10.1145/2766330.2766335>
39. Rweyemamu, W., Lauinger, T., Wilson, C., Robertson, W., Kirda, E.: Clustering and the weekend effect: Recommendations for the use of top domain lists in security research. In: Proc. of the Passive and Active Measurement Conference (PAM). pp. 161–177. Lecture Notes in Computer Science, Springer (2019)
40. Scheitle, Q., Hohlfeld, O., Gamba, J., Jelten, J., Zimmermann, T., Strowes, S.D., Vallina-Rodriguez, N.: A long way to the top: Significance, structure, and stability of Internet top lists. In: Proc. of the Internet Measurement Conference (IMC). pp. 478–493. ACM (2018)
41. Schinazi, D., Pauly, T.: Happy Eyeballs Version 2: Better Connectivity Using Concurrency. RFC 8305 (Proposed Standard) (Dec 2017), <https://www.rfc-editor.org/rfc/rfc8305.txt>
42. Sommesse, R., Moura, G.C.M., Jonker, M., van Rijswijk-Deij, R., Dainotti, A., Claffy, K.C., Sperotto, A.: When Parents and Children Disagree: Diving into DNS Delegation Inconsistency. In: Sperotto, A., Dainotti, A., Stiller, B. (eds.) Proc. of the 15th Passive and Active Measurement Conference (PAM). pp. 175–189. Lecture Notes in Computer Science, Springer (2020)
43. Streibelt, F., Böttger, J., Chatzis, N., Smaragdakis, G., Feldmann, A.: Exploring EDNS-Client-Subnet Adopters in Your Free Time. In: Proc. of the Internet Measurement Conference (IMC). p. 305–312. ACM (2013). <https://doi.org/10.1145/2504730.2504767>
44. ThousandEyes Blog, Cisco: The DDoS Attack on Dyn’s DNS Infrastructure. <https://www.thousandeyes.com/blog/dyn-dns-ddos-attack/>
45. Wing, D., Yourtchenko, A.: Happy Eyeballs: Success with Dual-Stack Hosts. RFC 6555 (Proposed Standard) (Apr 2012), <https://www.rfc-editor.org/rfc/rfc6555.txt>, obsoleted by RFC 8305

A DNS Resolution Overview

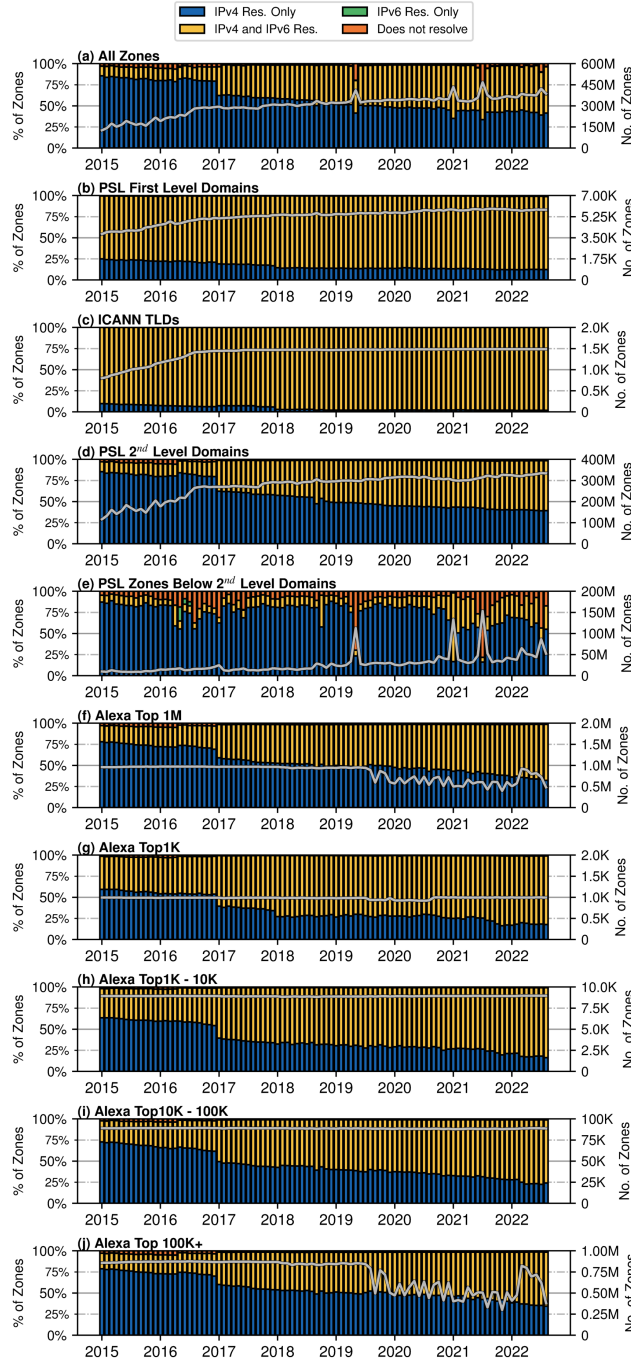


Fig. 6: Total number of zones in the dataset per month (gray line) and resolvability

B IPv6 Only Resolution Failures

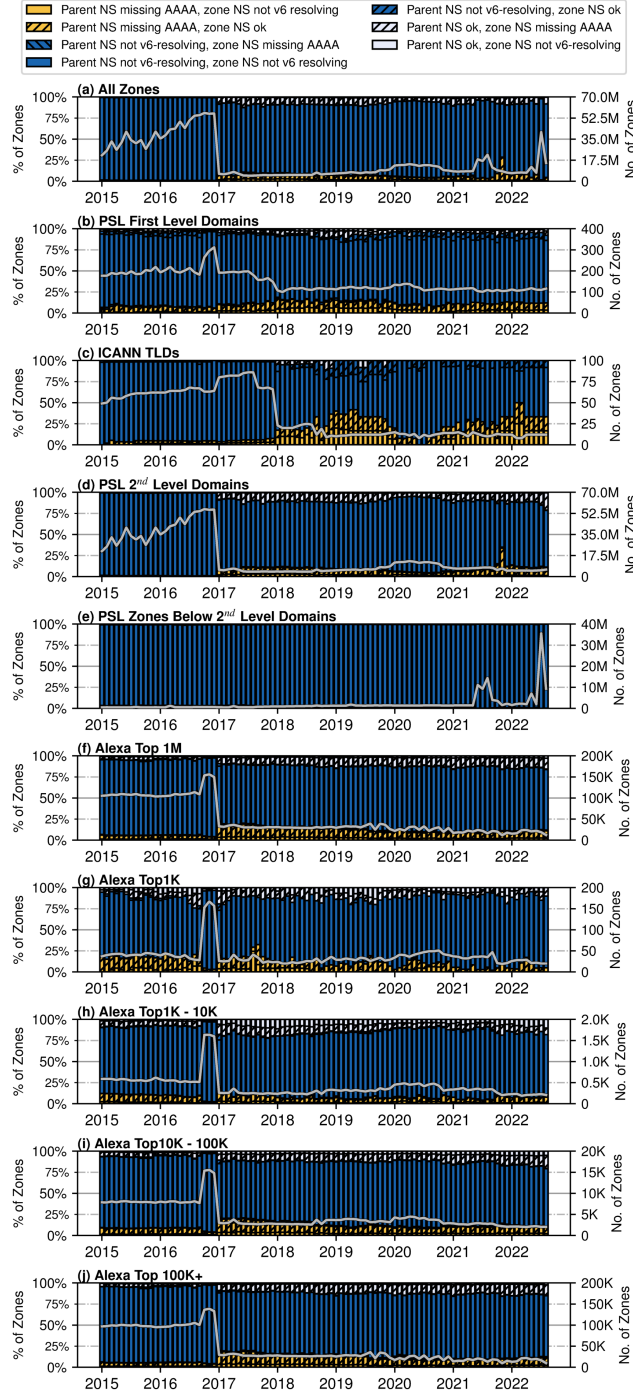


Fig. 7: Zones unable to resolve using IPv6, but with AAAA records in GLUE or zone apex (gray line), by resolution failure.

C Zones Without IPv6 Resolution per NS set

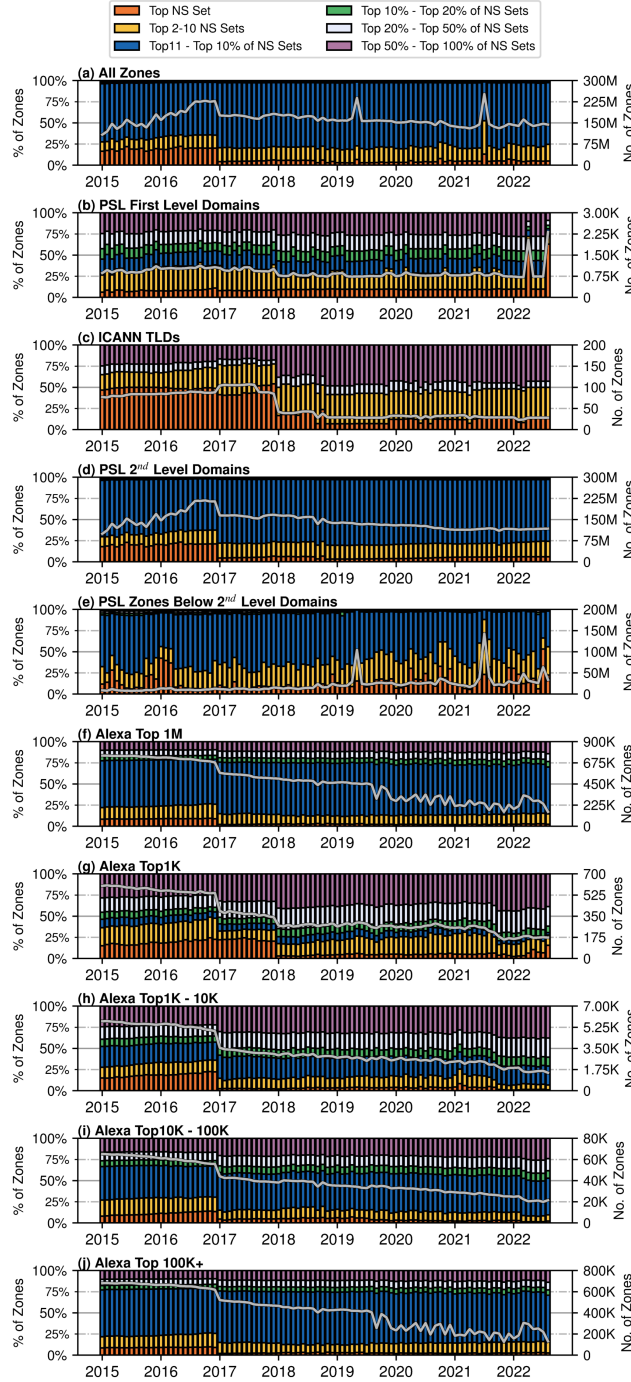


Fig. 8: Distribution of zones not resolving via IPv6 over NS sets.