A Technique for Counting DNSSEC Validators

Kensuke Fukuda
kensuke@nii.ac.jp
National Institute of Informatics

Shinta Sato
shinta@jprs.co.jp
Japan Registry Services

Takeshi Mitamura
mitamura@jprs.co.jp
Japan Registry Services

***Abstract*—The DNS security extensions (DNSSEC) is a new feature of DNS that provides an authentication mechanism that is now being deployed worldwide. However, we do not have enough knowledge about the deployment status of DNSSEC in the wild due to the difficulty of identifying DNSSEC validators (caching validating resolvers). In this paper, a simple and robust method is proposed that estimates DNSSEC validators from DNS query data passively measured at the server side. The key idea of the estimation method relies on the query patterns of the original query and the DNSSEC queries triggered by the original query, which is the ratio of the number of DS queries to the number of total queries per host (DSR: DS ratio). To show the effectiveness of the proposed method, we analyze passive traffic traces measured for all the ".jp" servers and actively send DNSSEC validation requests to caching resolvers that appear in the traces to obtain the ground truth data of DNSSEC validators. Our results of the active measurement reveal that less than 50% of the potential DNSSEC validators were validating caching resolvers in the wild; the remainder was related to stub validators (e.g., browser plugins) behind non-validating caching resolvers. Thus, simple IP address-based counts overestimated the number of DNSSEC validators in an investigation of the deployment of DNSSEC at the organization level (e.g., ISPs). Then, we demonstrate the effectiveness of the DSR by using the active and passive traffic traces. In summary, the ratio of validating caching resolvers in our dataset was estimated to be approximately 70% of the potential DNSSEC validators, and also 15-20% of the ASes sending DNSSEC queries were overestimated as ones with validating caching resolvers. In particular, our results show that some ASes providing public DNS service had few validating caching resolvers though they had a large number of hosts sending DNSSEC queries.**

I. Introduction

The Domain Name System (DNS) is one of the most important components of the current Internet. However, the original DNS had no functionality for authenticating queries between cache resolvers and authoritative servers against threats (e.g., cache poisoning [5]). The DNS security extensions (DNSSEC) [2], [4], [3] is a new feature of the DNS that verifies the correctness of query responses on the basis of the concept of a "chain of trust" from the trust anchor for the root zone to the edge authoritative zone entry. Two new resource records (RRs) are mainly introduced in DNSSEC: the DNS Key (DNSKEY) and Delegation Signer (DS). The DNSKEY RR of a zone entry is a public key for the zone entry stored on an authoritative server responsible for the zone, and the DS RR of the zone entry is a signature corresponding to the DNSKEY RR stored on the upper level authoritative server used for authorizing the corresponding DNSKEY. A DNSSEC validator requires pairs of DS and DNSKEY RRs for all corresponding zones in order to validate the correctness of the original query response.

The DNSSEC mechanism is still now being deployed all over the world. The ICANN signed the root zone and published a trust anchor in October 2010; subsequently, for example, the DS RR of ".jp" was registered on root servers in December 2010. However, we still do not have enough information about the deployment of DNSSEC validators in the wild though increases in the volume of DNS traffic have been reported [19]. One of the main difficulties comes from the noisy nature of DNS queries [22], [8]. Another reason is due to two different types of validatiors: validating caching resolvers (i.e., cache resolvers) and stub validators (e.g., a web browser plugin). The number of DNSSEC users is determined mainly by the deployment of validating caching resolvers because such resolvers have many clients, especially in large organizations. Thus, understanding the deployment status of the two different usages is essential for sketching out a DNSSEC deployment plan.

In this paper, we propose a simple and robust method that identifies validating caching resolvers in DNS traffic traces passively measured at authoritative servers. The key idea of the estimation is to use the ratio of the number of DS queries to the number of total queries per DNS host. We investigate the effectiveness of the proposed method by using DNS query logs passively collected at all name servers of ".jp" (the JP servers). Additionally, information on DNSSEC-available open resolvers appearing in the passive traces was collected as the ground truth data of the DNSSEC validators by actively sending the DNSSEC validation requests.

Our study has following four findings: (1) Only 50% of open resolvers that sent DNSSEC queries were DNSSEC validators; the remainder of them was simple cache resolvers, meaning that simple host counts overestimate the number of DNSSEC validators at the organization level. (2) The proposed heuristics, DS ratio, is a simple and robust index for use in distinguishing validating cache resolvers from non-validating cache ones. (3) We estimated that 70% of the DNS resolvers sending DNSSEC queries were validating caching resolvers in our dataset. (4) While the ratio of validating caching resolvers is stable in most ASes, some ASes providing public DNS servers had many non-validating caching resolvers.

II. Preliminaries

A. Related work

Many studies have been devoted to characterizing the behavior of DNS queries at authoritative servers, especially

root servers [9], [7], [22], [18], [8]. A common difficulty in characterizing DNS traffic is its noisy nature. For example, 90% of queries at the root servers have been found to be invalid [22], [8]. Additionally, the highly distributed nature of the DNS makes passive measurements difficult in terms of the coverage of the dataset though all queries first arrive at a root server.

As a live survey of DNS, Ref.[20] actively measured and analyzed a wide variety of statistics related to DNS cache servers on the Internet. One interesting finding is related to the popularity of DNS software. As of 2010, 34% of DNS servers used BIND software [13], the highest percentage for any software though the software for 40% of the servers was not identified. Many aspects of the DNS have also been analyzed such as the cache hit ratios [14], the availability of cache resolvers [17], and the performance differences between local and third party resolvers [1] in the measurements of caching resolvers. Furthermore, DNS query information has been shown to be useful in identifying malicious domains [23], [12], [6], [11], an important aspect of spam and botnet detection.

More specifically, in regard to the traffic behavior of DNSSEC, Ref. [19] investigated the impact of deploying DNSSEC on root servers, and found significant increases in the volume of DNS traffic due to DNSSEC deployment. Ref. [20] reported the number of DNSSEC available zones and the validation status for such zones by using active measurements. Similarly, Ref. [16] investigated the deployment and availability of DNSSEC with active measurements. In terms of the identification of DNSSEC hosts, Ref. [15] proposed criteria for use in identifying DNSSEC hosts. The key aspect of their criteria was to check the periodic arrival of DNSKEY RRs and their default TTL value. Furthermore, observing several types of failures in DNSSEC (i.e., software bugs, mismatches between the DS and DNSKEY, expiration of RRs) in the wild, Ref. [10] analyzed the impact of the validation failure at authoritative servers on validating caching resolvers.

B. DNSSEC validation

The key to authenticating the "chain of trust" of DNSSEC validation is the use of the two new resource records. As defined above, a DNSKEY RR is a public key for zone entry stored on the authoritative server responsible for the zone, and a DS RR is a signature corresponding to a DNSKEY RR stored on the upper-level zone used for authorizing the corresponding DNSKEY. The validator collects all pairs of DNSKEY and DS RRs from the root server to the authoritative server responsible for the zone entry, and it then validates all of them in a chain of trust manner. The validation fails if the DS or DNSKEY is not registered at an authoritative server or if one of the RRs is corrupted. Thus, the validator needs appropriate pairs of DS and DNSKEY RRs for validation.

As mentioned, validators are categorized into two types depending on the location of validation. The first is a cache resolver capable of DNSSEC validation (*validating caching resolver*) that does not need end hosts behind it to carry out validation. The other is a *stub validator* in which application software like that for a Web browser independently validates queries on an end host on its own. In this case, a conventional caching resolver does not validate queries from a stub validator (non-validating caching resolver), but does cache. Thus, with the second type, the client explicitly requests all DNSSEC-related queries by its own. However, distinguishing the two validation locations at an authoritative server is difficult because non-validating caching resolvers transfer DNSSEC queries from clients to authoritative servers even though they have no validation mechanism of their own.

C. Dataset

TABLE I
TRAFFIC BREAKDOWN

	Total		DNSSEC	
Data	IP addrs	ASes	IP addrs	ASes
201106	2150958	28722	9629	1705
201112	2081826	29600	14085	2490
201204	1904610	29334	21238	3295

Characterization of DNSSEC validators requires the collection of all queries related to DNSSEC as well as the original query. Due to the highly distributed nature of authoritative server deployment in each zone, it is sometimes difficult to measure all queries at authoritative servers (i.e., the root servers or ccTLD servers). Moreover, the noisy nature of DNS queries makes analysis more difficult. The data set we analyzed was composed of three 48-hour packet traces measured for all seven name servers of ".jp" (the JP servers) in Jun. 2011, Dec. 2011, and Apr. 2012 (referred to respectively as 201106, 201112 and 201204). The JP servers consist of seven servers distributed over the world, and most of them support anycast and replication. We focused mainly on the queries of the two RRs registered at the JP servers. The first were queries for "DNSKEY .jp" and the second were DS queries for a subdomain of ".jp" (e.g., "DS example.jp"). DNSSEC validation requires both RRs from the JP servers. The TTL value for the DNSKEY and DS in .jp was 24 h; thus we expected to find at least one query for "DNSKEY .jp" from a valid and busy DNSSEC host depending on the query timing. Given the completeness of our dataset, all the queries to and from the JP servers were basically captured in our measurements. This enabled us to investigate the total behavior of cache resolvers related to the JP servers, unlike a previous study [15] while the coverage of our data was limited to JP domains. The breakdown of the IP addresses and ASes for all queries and DNSSEC queries is shown in Table I. DNS resolvers sending at least one DNSSEC query are regarded as potential DNSSEC validators.

Furthermore, to obtain the ground truth data set for the validating caching resolvers, we actively sent a query turning on the DNSSEC OK bit (DO), which requests DNSSEC validation, to all IP addresses that sent at least one DNSSEC query in the passive traces. The reply message from a host

should contain an AD flag if it has validated a query successfully, meaning that it is a validating caching resolver. A *non-validating caching resolver* replies to this without an AD flag, on the other hand, meaning that it is a simple cache resolver with stub validators (or other validating caching resolvers) as clients.

III. Dynamics of DNSSEC validators

A. Actual validating caching resolvers

TABLE II
VALIDATING AND NON-VALIDATING CACHING RESOLVERS

201106	validator	non-validator	total
DS-DNSKEY	276 (52%)	257 (48%)	533
DS-only	48 (18%)	226 (82%)	274
DNSKEY-only	7 (11%)	63 (89%)	70
total	331 (38%)	546 (62%)	887

201112	validator	non-validator	total
DS-DNSKEY	546 (61%)	343 (39%)	889
DS-only	109 (34%)	208 (66%)	317
DNSKEY-only	26 (16%)	133 (84%)	159
total	681 (50%)	684 (50%)	1365

201204	validator	non-validator	total
DS-DNSKEY	891 (64%)	510 (36%)	1401
DS-only	125 (17%)	602 (83%)	727
DNSKEY-only	13 (11%)	109 (89%)	122
total	1029 (46%)	1221 (54%)	2250

We first examined the actual validating and non-validating cache resolvers in the traces. As the ground truth data, we used the IP addresses in passive traces that replied to active probes with DO bit $= 1$ (i.e., open resolvers). The open resolvers replying to these probes with AD flags were defined as validating caching resolvers, and those replying without them were defined as non-validating caching resolvers. Note that stub validators do not appear as open resolvers because they simply ignored the probes. The total numbers of open resolvers sending DNSSEC queries increased over the months; 877 for 201106, 1365 for 201112, and 2250 for 201204, which consistently accounted for about 10% of the potential DNSSEC validators summarized in Table II. The percentages of validating caching resolvers were stable (38, 50, and 46%, respectively). Thus, more than half of the open resolvers that appeared in the traces did not validate the queries on their own; they simply forwarded the queries between stub validators and authoritative servers acting as cache resolvers.

We further categorized the potential DNSSEC validators into three types for investigating the validity of DNSSEC validators: *DS-DNSKEY*, a host sending both DS and DNSKEY queries, *DS-only*, a host sending only DS queries, and *DNSKEY-only*, a host sending only DNSKEY queries. The percentages of validating caching resolvers were higher for DS-DNSKEY hosts (52, 61, and 64%, respectively) because they were more plausible validators. We emphasize that a non-negligible number of DS-only hosts validated DNSSEC queries as actual validating caching resolvers even though DNSKEY-only hosts had less possibility of being actual validating caching resolvers. In summary, these results indicate the difficulty of identifying DNSSEC validators by using the simple appearance of DNSKEY and/or DS queries even if all the queries are captured.

B. Characterizing validating caching resolvers

Next, we focused on the relationship between the number of DS queries and the number of other queries per DNS resolver. We assumed that the number of DS queries per validator has a positive correlation with that of other queries because DS queries are triggered by other queries. Similarly, a non-validating caching resolver was assumed to receive a small fraction of the DS queries from a small number of validators compared with the other queries from a large number of non-validators. We demonstrate the scatter plots for the number of all queries and those of the DS queries per DNS resolver in Fig. 1, which shows plots for (a) the validating caching resolver and DS-DNSKEY, (b) the validating caching resolver and DS-only, (c) the non-validating caching resolver and DS-DNSKEY, and (d) the non-validating caching resolver and DS-only. The plots for the validators ((a) and (b)) are clearly concentrated along the diagonal, showing a positive correlation. Thus, the majority of validators explicitly sent the DS queries as expected. In comparison, the behavior of the non-validators ((c) and (d)) is more complicated than that of the validators. Some plots can be distinguished along the diagonal even though the rest of them are spread widely especially in the area corresponding to a larger number of all queries and a smaller number of DS queries. The former can be viewed as a stable and small non-validating caching resolver that has a few stub validators and a few end hosts (non-validator) as well as one that has a few validating caching validators functioning as a forwarder. Thus, this behavior resembles the behavior of a validating caching resolver that has fewer end hosts. In other words, it is hard to distinguish between a validating caching resolver and a non-validating caching resolver (stub validator) in this case. The latter is typical for a large non-validating caching resolver for which a few stub validators send DS queries hidden among a huge number of queries from other end hosts. Another key point here is that even for DS-only hosts, a diagonal-based representation is still useful for identifying validating caching resolvers as can be seen in (b) and (d).

IV. Estimating the number of validating caching resolvers

On the basis of above observations, we define a ratio (the DSR), which is the number of DS queries to that of all queries per host, and used it to characterize validators. DSR ranges from zero to one, and the lower the DSR (i.e., the greater the deviation from the diagonal), the greater the possibility of a host being a non-validating caching resolver. The goal of this section is to estimate the number of validating caching resolvers in the original traces with the DSR.

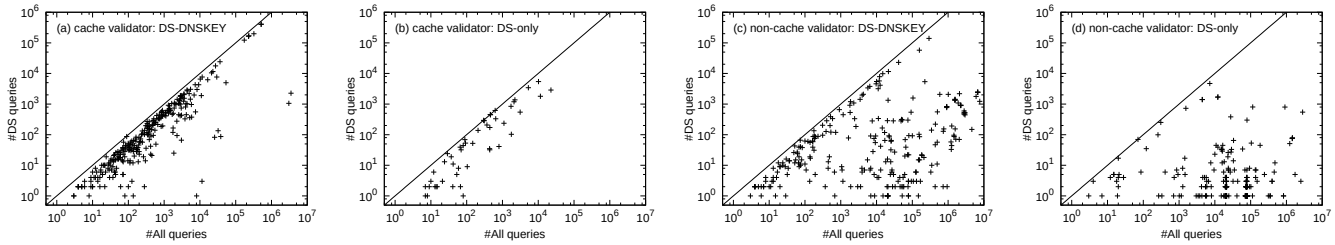

Fig. 1. Scatter plots of the number of all queries and the number of DS queries (2011.06) (a) validating caching resolvers (DS-DNSKEY), (b) validating caching resolvers (DS-only), (c) non-validating caching resolvers (DS-DNSKEY), and (d) non-validating caching resolvers (DS-only)

TABLE III
CONFUSION MATRIX BASED ON DSR

(201106)	validator	non-validator	total
DSR $>$ 0.04	307 (TP)	125 (FP)	432
DSR $\leq$ 0.04	17 (FN)	358 (TN)	375
total	324	483	807
(201112)	validator	non-validator	total
DSR $>$ 0.04	625 (TP)	150 (FP)	775
DSR $\leq$ 0.04	30 (FN)	401 (TN)	431
total	655	551	1206
(201204)	validator	non-validator	total
DSR $>$ 0.04	964 (TP)	265 (FP)	1229
DSR $\leq$ 0.04	52 (FN)	847 (TN)	899
total	1016	1112	2128

Table III lists a confusion matrix of validating and non-validating caching resolvers for the best threshold value of DSR = 0.04. We obtained 307 true positive (TP) and 358 true negative (TN) hosts for 201106. The number of false positives (FPs) was slightly higher (125 hosts) than that of false negatives (FNs) (17 hosts). The number of FPs is attributed to the hosts being distributed near the diagonal in Fig. 1 (c). In other words, the FP hosts were interpreted as stable stub validators. We used several commonly used performance indices for the classification: accuracy $= \frac{TP+TN}{TP+TN+FP+FN}$, precision $= \frac{TP}{TP+FP}$, recall $= \frac{TP}{TP+FN}$, and f-measure $= \frac{2precision \times recall}{precision+recall}$. These indices range between zero and one, and a higher value indicates better performance. As shown by the summarized results in Table IV, the DSR for the best threshold was an effective feature for classification.

TABLE IV
PERFORMANCE INDICES (DSR = 0.04)

	accuracy	precision	recall	f-measure
201106	0.82	0.71	0.95	0.81
201112	0.85	0.81	0.95	0.87
201204	0.85	0.78	0.95	0.86

We estimated the number of validators by using the DSR and the original three traces. As shown in Table V, the ratio of validating caching resolvers was consistent among the traces (70, 76, and 73%, respectively). Furthermore, we concluded that 15-20% of the ASes sending DNSSEC queries were overestimated as ASes with validating caching resolvers.

Finally, we plotted the number of potential DNSSEC validators and that of the estimated validating caching resolvers per AS for 201106 and 201204 in Fig. 2. Each plot corresponds to an AS, and plots below the horizontal dotted lines represent ASes without estimated validating caching resolvers. Comparison of the two sub figures showed that the number of ASes increased from 1705 to 3295 and that the number of DNSSEC hosts increased in many ASes. While most plots have a high correlation between the two numbers, we found that a number of estimated validating aching resolvers in some ASes is much smaller than that of the DNSSEC hosts. For example, labels (A) are ASes characterized by a number of DNSSEC hosts but no estimated validating caching resolvers though the number of potential DNSSEC validators increased in both ASes. We also confirmed that most of such potential DNSSEC validators in both ASes had no DNSKEY queries. These ASes provide public DNS servers without DNSSEC validation, so their cache resolvers likely send queries as non-validating cache resolvers. Similarly, ASes denoted by (B) (also providing a public DNS server) have fewer validating caching resolvers than the observed potential DNSSEC validators. In contrast, labels (C), which corresponded to some large ISPs, fall on the diagonal, meaning that their cache resolvers are more likely to perform validation.

TABLE V
NUMBER OF ESTIMATED VALIDATORS

(DS-DNSKEY & DS-only)	201106	201112	201204
DSR $>$ 0.04	5903 (70%)	9043 (76%)	13830 (73%)
DSR $\leq$ 0.04	2494 (30%)	2922 (24%)	5201(27%)
total	8397	11965	19031

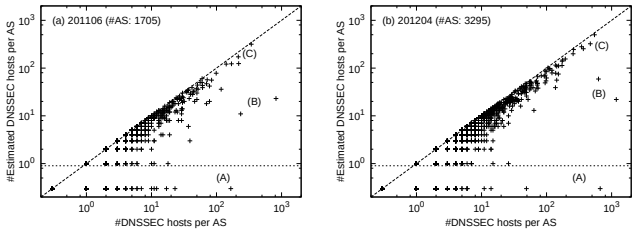

Fig. 2. Scatter plots of the number of potential DNSSEC validators and that of estimated validating caching resolvers per AS

V. Discussion

As we pointed out, counting the number of DNSSEC validators in the dataset is difficult. One of the reasons is shorter measurement period (48h) compared with the default TTL of RRs (24h). While checking the inter-arrival period of DNSKEY queries [15] is a promising approach to reliably estimating the number of DNSSEC validators, our results showed that a non-negligible number of DNSSEC validators sent zero or one DNSKEY query in the observed time period. Moreover, such hosts were shown by the ground truth data to actually be validating caching resolvers. To solve this problem, we showed that the DSR works well for classifying validating and non-validating caching resolvers. In addition, the DSR is robust against missing (or sampling) data. Compared with using pairs of queries [15], the cost of calculating the DSR is small, and it works even if some queries were missed. It accurately identified validating caching resolvers from among DS-only hosts. DSR could also work as a filter against noisy DNS queries. However, since our ground truth data relied on the behavior of the open resolvers, there might have been a bias that differed from that of other cache resolvers. Nevertheless, since the idea of DSR was derived from the basic query pattern of the DNSSEC, the effect of this bias should be small.

A potential error of the DSR may cause the host count to double. If queries from a stub validator are equally dispatched to multiple non-validating caching resolvers, their DSRs could be higher. One possible scenario for this will be a cache resolver with both IPv4 and IPv6 addresses. In addition, our results heavily depended on the current deployment of DNS software; de facto BIND [13], accounting for over 40% of the cache resolvers on the Internet [20], explicitly queries DS RRs without using these in additional fields while unbound [21] uses these fields explicitly, even though RFCs do not specify how DS RRs are treated in additional fields of the preceding queries.

VI. Conclusion

We analyzed emerging DNSSEC traffic and the behavior of validating caching resolvers by using passive and active measurements. Our active measurements of open resolvers revealed that less than 50% of potential DNSSEC validators were actually validating caching resolvers; the remainder was related to stub validators behind non-validating caching resolvers in the wild. Thus, simple IP-address-based estimations overestimate the number of DNSSEC validators in terms of the deployment of DNSSEC at the organization level (e.g., ISPs). We introduced a simple and robust traffic feature for use in distinguishing validating caching resolvers from non-validating ones that is based on the ratio of the number of DS queries to all queries and demonstrated its effectiveness by using ground truth data. Finally, we estimated that the ratio of validating caching resolvers in our dataset was approximately 70% of the potential DNSSEC validators and that 15-20% of the ASes sending DNSSEC queries were overestimated as ones with validating caching resolvers. In particular, some ASes providing public DNS service had few validating caching resolvers though they had a large number of hosts sending DNSSEC queries.

Acknowledgement

We are grateful to thank Kenjiro Cho and Romain Fontugne for their helpful discussion. We also thank the anonymous reviewers for their helpful comments. This work was partially supported by a Grant-in-Aid for Young Scientists (A), from the MEXT/JSPS in Japan.

References

[1] B. Ager, W. Muehlbauer, G. Smaragdakis, and S. Uhlig. Comparing DNS Resolvers in the Wild. In *IMC2010*, pages 15–21, Melbourne, Australia, Nov 2010.

[2] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. DNS Security Introduction and Requirements. RFC4033, Mar 2005.

[3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. Protocol Modifications for the DNS Security Extensions. RFC4035, Mar 2005.

[4] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. Resource Records for the DNS Security Extensions. RFC4034, Mar 2005.

[5] D. Atkins and R. Austein. Threat analysis of the domain name system. RFC3833, Aug 2004.

[6] L. Bilge, E. Kirda, C. Kruegel, and M. Balduzzi. EXPOSURE: Finding Malicious Domains Using Passive DNS Analysis. In *NDSS2011*, San Diego, CA, Feb 2011.

[7] N. Brownlee, kc Claffy, and E. Nemeth. DNS Root/gTLD Performance Measurements. In *USENIX LISA*, San Diego, CA, Dec 2001.

[8] S. Castro, D. Wessels, M. Fomenkov, and kc Claffy. A Day at the Root of the Internet. *Computer Communication Review*, 38(5):41–46, Oct 2008.

[9] P. B. Danzig, K. Obraczka, and A. Kumar. An Analysis of Wide-Area Name Server Traffic: A Study of the Internet Domain Name System. In *SIGCOMM1992*, pages 281–292, Baltimore, MD, Aug 1992.

[10] K. Fukuda, S. Sato, and T. Mitamura. Preliminary evaluation of potential impact of failure in dnssec validation. In *DNSEASY2011*, page 12, Rome, Italy, Oct 2011.

[11] S. Hao, N. Feamster, and R. Pandrangi. Monitoring the initial dns behavior of malicious domains. In *IMC2011*, pages 269–278, Berlin, Germany, Nov 2011.

[12] T. Holz, C. Gorecki, K. Rieck, and F. C. Freiling. Measuring and detecting fast-flux service networks. In *NDSS2008*, San Diego, CA, Feb 2008.

[13] Internet Systems Consortium. BIND. http://www.isc.org/software/bind.

[14] J. Jung, E. Sit, H. Balakrishnan, and R. Morris. DNS Performance and the Effectiveness of Caching. *IEEE/ACM Transactions on Networking*, 10(5):589–603, Oct 2002.

[15] Ólafur Gusmundsson and S. D. Crocker. Observing DNSSEC validation in the wild. In *Workshop on Securing and Trusting Internet Names (SATIN2011)*, page 8, Teddington, UK, Apr 2011.

[16] E. Osterweil, M. Ryan, D. Massey, and L. Zhang. Quantifying the Operational Status of DNSSEC Deployment. In *IMC2008*, pages 231–241, Vouliagmeni, Greece, Oct 2008.

[17] J. Pang, A. Akella, J. Hendricks, R. D. Prisco, B. Maggs, and S. Seshan. Availability, Usage, and Deployment Characteristics of the Domain Name System. In *IMC2004*, pages 1–14, Sicily, Italy, Oct 2004.

[18] Y. Sekiya, K. Cho, A. Kato, R. Somegawa, T. Jinmei, and J. Murai. Root and ccTLD DNS server observation from worldwide locations. In *PAM2003*, La Jolla, CA, Apr 2003.

[19] G. Sisson. Deployment of DNSSEC in the Root Zone: Impact Analysis. Technical report, DNS-OARC, Dec 2010. https://www.dns-oarc.net/files/DURZ_Report_Final.pdf.

[20] G. Sisson. DNS Survey: October 2010. http://dns.measurement-factory.com/surveys/201010/, 2010.

[21] Unbound. http://unbound.net/.

[22] D. Wessels and M. Fomenkov. Wow, That's a Lot of Packets. In *PAM2003*, La Jolla, CA, Apr 2003.

[23] B. Zdrnja, N. Brownlee, and D. Wessels. Passive Monitoring of DNS Anomalies. In *DIMVA2007*, pages 129–139, Lucerne, Switzerland, Jul 2007.