Observing DNSSEC validation in the wild

Ólafur Guðmundsson and Stephen D. Crocker
Shinkuro, Inc.
4922 Fairmont Avenue
Suite 250
Bethesda MD 20814 U.S.A.
Email: ogud@shinkuro.com and steve@shinkuro.com

***Abstract*—DNSSEC protocol deployment has taken place in phases, beginning with protocol development and followed by the signing of top-level zones and early-adopter "leaf" zones. The next phase is to encourage wide-scale validation, as that will improve the overall DNS system and enable new applications. In order to quantify DNSSEC usage for audiences it is important to be able to measure how many zones are signed and how widespread validation is. This paper will describe how to measure validation by looking at DNS queries; in it, we present results from two sample periods monitoring a sub-set of the authoritative name servers for .org.**

***Index Terms*—DNSSEC, DNS operations, Internet measurements.**

I. Introduction

In this paper we will discuss how to measure the population of resolvers performing DNSSEC[1], [2], [3] validation. Knowing the prevalence of validation among resolver populations has a number of possible uses, including measurement of global trends, identification of communities that generally use DNSSEC validation, auditing, etc.

A. Different types of resolvers from the DNSSEC point of view

The standard classification of resolvers into recursive and stub resolvers does not apply to DNSSEC validation. What we are interested in is what kind of DNSSEC behavior the resolvers show. We classify resolvers into three groups:

A) Validating resolvers
B) DNSSEC-capable but not validating
C) DNSSEC-ignorant

In this discussion we generally make the following simplification: resolver = IP address. In many cases, recursive resolvers are located behind a NAT device and we see evidence that there are frequently multiple resolvers behind a single IP address. When there are multiple resolvers, we classify the address by highest observed behavior (of items A-C above), and classify the resolvers by looking at the content of a DNS query and DNSSEC-specific query patterns.

DNSSEC-ignorant resolvers (C above) never ask a question with the EDNS0 and DO bit set, so if all questions seen from an address feature this behavior, we classify the resolver as DNSSEC-ignorant. This is not strictly correct, however, as there are DNS proxies that strip EDNS0, and when multiple resolvers are behind a NAT then we cannot tell the differences between each one.

B. Goals of this work: Be able to monitor DNSSEC validation growth

We want to develop simple rules and techniques that can be built into DNS monitoring systems and can issue cumulative results all the time, not just periodically as in the past. If that is not possible, we seek to define a subset of the data that DNS operators collect, which will allow validation measurements. Large DNS operators have to be able to handle enormous volumes of DNS traffic, both to answer queries and frequently to capture it for incident post-mortems.

II. Validation of Previous Work

In the past, observers have monitored how many zones are signed and examined other evidence of DNSSEC deployment[4]. They have also looked at DNS cache performance[5], how different DNS resolver operators perform[6], etc. In addition, several studies and extensive monitoring by root operators was performed in the last few years to measure the deployment of EDNS0 and buffer sizes advertised[7], [8], and the impact of signing the root zone.

These studies are not directly related to our work, but in the following discussion we build on them.

III. Observing Validating Resolvers

DNSSEC[3] extends the old DNS protocol by signing RRsets, which provides source authentication. In order to perform validation, a chain of trust must be built from a trust anchor or the root of the DNS tree. The trust chain consists of two new RR types:

- DNSKEY, which contains a public part of a key. This record type is stored at the apex (top of) the zone the key will sign.
- DS is a cryptographic hash of the binding of a domain name and the public key used to sign the DNSKEY set at the domain. The DS record is stored at the parent and is handed out as part of the referral returned by the parent when it receives a request for a domain name below a zone cut in one of the zones it is authoritative for.

A. Query Patterns to Look For

Validating resolvers have slightly different query patterns from non-validating ones, and will explicitly ask for records in the DNSSEC trust chain (DNSKEY and DS) that are almost never asked for by non-validating resolvers. We will explain below how each one works.

1) Monitoring DO bit : Resolvers that understand DNSSEC record types usually set the DO (DNSSEC OK)[3] bit on queries. The presence of this bit can be seen as indicating that the querier can validate DNSSEC answers. Thus this resolver can be in category A or B as described in Section 1.A above.

2) Monitoring DNSKEY queries: Every time a validating resolver attempts to validate information from a zone (secure.example.) for the first time, it will get from the parent authoritative server (example.) a referral that contains a DS RRset. It will then ask secure.example's authoritative server for the resource record it seeks, followed by a query for the DNSKEY record so that it may validate the resource record in the prior query.

After the DNSKEY record has timed out of the validating resolvers cache, it needs to re-fetch the DNSKEY record in order to validate any future answers from the zone.

3) Monitoring DS queries: A parent zone can monitor validation performed against a child zone by looking at DS queries. Each time the DNSKEY from the child times out (and if the DS record has timed out), the validating resolver needs to fetch both the DS record from the parent (example.) and the DNSKEY RRset from the child (secure.example). When a DNSKEY RRset expires from the cache in a validating resolver, the DNSKEY set needs to be fetched and validated the next time this set is needed to validate a new RRset. If the corresponding DS RRset has also expired, the DS set needs to be fetched and validated.

B. Rules used for determining resolver category

In order to process the traces we examine well set of rules identify validating resolvers and if different rules support each other's results. The rules we applied are:

A.1 DS query, definitly a validating resolver (in fact each DS query increases the probability it is validating, but for now we are assuming one query means validation)

A.2 DNSKEY followed by DNSKEY after TTL, this implies a validating resolver

A.3 Query for X.zone followed by DNSKEY right after, definitly a validating resolver

A.4 DNSKEY query seen. In this case we label the resolver as possibly validating,

We realize that we may be classifying more than just recursive resolvers; for example, there are monitoring tools, test queries and trust-anchor maintainers that may issue similar query patterns. We have decided that usage of these tools is closely related to DNSSEC deployment and will count them in this discussion.

While these rules may seem straightforward to apply there are complications depending on many factors including:

- Query scattering, this refers to how resolvers send questions to different authoratative name servers for successive queries
- Query ordering, when a resolver needs to look up multiple RRsets how are these queries issued
- Packet losses, if the resolver for some reason does not get an answer to a question it will repeat the question, this repeated question may or may not go to the same nameserver. In particular resolvers that advertise large EDNS0 buffer size but are behind a link that does not allow answer as large as the .org DNSKEY (just over 1500 bytes) at the time of measurements, in this case we are likely to see DNSKEY repeated in a short time.
- Multiple resolvers behind a NAT, these devices may send DNSKEY queries multiple times inside a DNSKEY TTL but spread out far enough that this is not a retry due to packet loss.
- Elimination of repeated queries.

If a complete set of traces over a long time from all authoritative name servers can be analyzed, it will be possible to build a good picture of how many of the resolvers in the domain's "working set" are validating.

C. Traffic selection

We seek to determine and measure worldwide DNSSEC validation and sought to analyze a domain used all over the world (and not just by the technical community). This led us to approach the the operator of the .org top-level domain (TLD), and after discussion with PIR (the .org "owner") and Afilias (the domain operator), we were in mid-2010 granted access to some traces from the name servers that Afilias operates for .org. This is only a partial set of name servers since processing the large volume of available data is challenging. As a result, we began by analyzing short (as in time) samples. Below are our initial results from analyzing this traffic and how, in general, we have attempted to measure DNSSEC validation in the wild.

1) Which DNSKEY query to look at? : We originally considered monitoring the explicit queries for .org's DNSKEY, then considered expanding our measurements to all DNSKEY queries for domains inside .org since this allows us to detect more validators and possibly overcome the scatter of DNS queries. Two domains inside .org seem to have significant DNSKEY traffic: isc.org and ietf.org. Afilias operates most (five out of six) of the name servers for ietf.org, so we have access to ietf.org's traffic. This indicates that DNSSEC validation is taking place for ietf.org, but traffic volume was not large enough for us to draw conclusions. isc.org has a large amount of DNSKEY traffic, a phenomenon that seems driven in large part by traffic related to dlv.isc.org, and by the fact that the DNSKEY query for isc.org seems to be used in DDoS attacks. DLV is a tool through which early DNSSEC adopters could validate via an alternative trust chain. We considered including dlv.isc.org look-ups in our statistics but decided against it, even though doing so would inflate the numbers of DNSSEC adopters, because we wanted only to measure the traffic by validators that validate to the root. Many of the resolvers that are looking for dlv.isc.org do not seem to have been configured with the root key. Investigating dlv.isc.org traffic in more detail should be interesting since, at first glance, it appears that there are more validators resolving

to the ISC DLV trust anchor than to the root at the time of measurements, although some are using both the root key and the DLV key.

2) Which DS query to look at? : We have counted all DS queries, which may include DS queries like "child.parent.org", org will only have a referral for "parent.org", but we decided that just seeing a DS query is enough evidence of validation that it was acceptable to count this resolver as validating.

D. Monitoring .org domain traffic

The main purpose of our work is to develop methodologies and tools to measure DNSSEC validation in the wild. .org was rare among TLDs in that its DNSKEY RR TTL is quite short—only 15 minutes, where most other TLDs have much longer TTLs for DNSKEY records. See Table 1 for some of the DNSKEY TTL and DS TTL values used in signed TLDs.

TABLE I
SAMPLE OF SIGNED TLD TTL VALUES

domain	DNSKEY TTL	DS TTL	NSEC TTL
. (root)	2 days	1 day	1 day
org	15 mins.	1 day	1 day
br	6 hours	1 day	15 mins.
net	1 day	1 day	1 day
se	1 hour	1 hour	2 hours
us	1 day	2 hours	1 day
fr	2 days	not found	90 mins.
cz	1 hour	5 hours	15 mins.
jp	1 day	1 day	15 mins.

At this point there is no agreement among operators or protocol experts about what constitute "good" values. A comprehensive timing analysis is needed to identify "better" values, but such an analysis is also frequently affected by the different goals of each operator. That topic is outside the scope of this paper.

We have periodically collected traces for short periods (30-50 minutes) to analyze and improve our tools and techniques.

1) .org name-server locations impact on query patterns : We have access to traces for the .org-registry name servers that Afilias operates. The name-server set for .org has six name servers and all addresses are any-cast. Four of the addresses are provided by Afilias while two are provided by an outside contractor, PCH. Afilias operates five sites around the world and PCH provides .org resolution in about 15 locations. The .Org nameservers are all any-cast servers; and highly distributed aound the world, we have traces from sites in North America, Europe and Asia. PCH provides sites on all continents.

The first question is how much of .org's traffic ends up on servers we have traces from. One might infer four-sixths or two-thirds, but that is not how DNS resolvers work. Most recursive DNS resolvers try to go to the "closest" authoritative server for a given domain, leading to two different behaviors for resolvers, scattering queries in the beginning while trying to figure out which authoritative server is closest. Once that has been established, most queries go to the closest one, with sporadic queries to the next-nearest server to check whether the distance to it has changed. Different resolvers have different criteria for selecting the "closest" authoritative server; for example, Unbound says all servers within [min_rtt .. min_rtt + 400ms] are equal, and with any-casting almost all the .org name servers will fall into this band. On the other hand, Bind-9.6 and Bind-9.7 by default have a much smaller band and thus will discriminate more among authoritative servers.[1]

For this reason, some busy resolvers will concentrate queries to a single address. But resolvers that do not send lots of queries will frequently scatter them, except the ones that do not care about performance and only send queries to the first name server in the set the resolver received. This behavior has three implications:

- Resolvers that resolve lots of .org names will either be in our sample or not at all.
- Resolvers that only do occasional .org queries will show up in our samples but we may not see enough of their queries to determine whether they are performing DNSSEC validation.
- Resolvers that do not do Round Trip Time (RTT) optimization may lock in on a "random" server.

The actual situation is that Afilias's servers answer about 50 percent of queries for .org. Based on this we need a model that tells us how to model what we see in the sample we have toward how the entire resolver population behaves. If there are geographical differences in DNSSEC validation, depending on where it is it may show up in the traces we analyzed. For example, if Japan has high concentration of DNSSEC validation we are unlikely to see much of it as PCH provides a servers there but the closest Afilias servers are in Singapore and Seattle.

2) Sample period and what to look for : One of the main reasons .org is a good domain to use to measure DNSSEC validation in the wild is the short TTL on its DNSKEY, which forces the validating resolver to re-query for the DNSKEY every 15-plus minutes. Of course the validator only fetches DNSKEY when it needs it, thus the 15 "plus." The DS TTL of one day, on the other hand, means that DS queries are not as common as in TLDs with short DS TTL. Many of the DS queries are for domains that have a different TTL on their NS and DNSKEY RRsets.

For example, in .br domains the DNSKEY TTL is 24 hours and the TTL on the NS record in the .org zone apex is also 24 hours, but the TTL on the DS records is six hours. Thus for domains that have a TTL on their DNSKEY of not equal to 24 hours, there is the potential for a DS query.

The sample period for .org has to be long enough to make it probable that a DNSKEY query happens multiple times for a busy validating resolver. For a non-busy validator, time of day may have a significant impact on results. One of our assumptions was that validator = IP address, but for longer

[1]Unbound scatters queries more, which makes DNSSEC validation easier to spot the scattering is done for cache-spoofing prevention. Bind is more tuned for performance. One of the big unknowns is what is the market shares of different kinds of resolvers. These numbers are unknown, but Bind-9 is most likely the largest base of resolvers.

periods that assumption may not hold as well and we would expect to see either multiple resolvers at a single address or a mobile resolver that shows up on multiple addresses.

Traces come in 10-minute increments. We originally selected 30 minutes as a sample size, but after some experience increased that to 50 minutes to increase the possibility of seeing multiple DNSKEY questions.

IV. Results

Due to the scatter of queries and the fact that we only see a subset of authoritative servers, we were unable to apply rule A.3. We have analyzed two sets of traces, one from early November 2010 and the second from early January 2011. Both traces include all of the Afilias sites, and both are from the same time of day, 2000-2050 UTC.

This time was chosen randomly and we expect that if there are significant differences in the geographical distribution of DNSSEC validators, the time of day chosen will have some impact on the results. What we are mainly interested in at this time is to see if there were changes in DNSSEC validation between the two periods..

The traces are full packet captures and thus quite large, requiring long transfer and processing times. We compacted the results to focus on the models above, although there are other significant topics to research in these traces such as TCP usage and attempts to understand client behaviors.

A. November 2010 results

Table 2 shows how many different resolvers asked for DS records from each site in each 10-minute interval. Here, "Seen" is the number of validators this state has seen and "Population" is the cumulative number of validators seen from the beginning by all sites.

TABLE II
DS queries seen in each period in November 2010

Interval	Site A	B	C	D	E	Seen	Population
0-10	1971	921	218	881	348	3066	3066
10-20	478	880	1146	823	342	2609	4028
20-30	1753	850	1124	832	297	3191	4765
30-40	1405	880	835	825	338	2816	5202
40-50	289	828	929	808	361	2315	5439
Total	2869	1894	1924	1682	651		

Table 3 shows how many resolvers passed the two DNSKEY look-up tests (see A.2 in Section 3.2 above). The large difference between Questions and Sources is to a large extent caused by validators that advertise a larger EDNS0 value than fits on their "site" link, as well as the fact that the .org DNSKEY set is over 1500 bytes and that many sites are known to not allow fragmented UDP packets through the site's firewall.

Table 3 shows that about 29 percent of the resolvers that issue DNSKEY queries pass our test but only 863 are certified by a single site, meaning that at least 312 validators pass because we are seeing questions at multiple sites. This is partially caused by the scattering of DNSKEY queries; even though we had 50-minute traces, the probability that we see a second DNSKEY query is still constrained by the scattering of queries and the validator's need to fetch the .org DNSKEY.

TABLE III
DNSKEY queries seen and results of DNSKEY criteria

	Questions	Sources	Passed	Pass %
A	3804	1337	118	8.8
B	7148	1169	238	20.4
C	3201	1097	179	16.3
D	3565	929	210	22.6
E	912	481	118	24.5
Total	18630	4045	1175	29.0

In addition, 993 validators are confirmed by both rules, and we have a total of 5621 validators confirmed, meaning that only 182 are confirmed by DNSKEY look-ups. Further, of the 2870 possible validators, 2764 are on the DS list, leaving only 106 possible unconfirmed validators.

B. January 2011 results

Table 4 shows the results for DS queries. It's clear that there are fewer DS sources than in the prior sample.

TABLE IV
January 2011 DS queries and sources

	A	B	C	D	E	Seen	Population
0-10	1469	681	641	614	350	2607	2607
10-20	796	533	725	457	340	1951	3209
20-30	707	711	375	604	343	1879	3453
30-40	1384	406	320	473	307	2140	3817
40-50	782	345	259	593	320	1630	3970
Total	2171	1171	1034	1149	686	3970	

Table 5 shows the DNSKEY look-up test results. Adding up the Passed column for all sites in Table 4 yields 615, meaning there are at least 143 validators that passed only by correlating multiple sites. The strong tendency of validators to successfully pass this test at a single site indicates that there are many more that we cannot see, and we need to build models allowing us to estimate their numbers.

TABLE V
January 2011 DNSKEY

	Questions	Sources	Passed A.2 test	Pass %
A	3889	1519	248	16.3
B	3364	859	99	11.5
C	1316	496	36	7.3
D	1795	749	134	17.9
E	964	511	98	19.2
Total	11196	3348	758	22.6

Of the 2590 possible validators from this result, 1805 are on the DS list, leaving 785 as possible validators. This gives us 4728 confirmed validators.

C. Comparing the two sets' results

At first glance, Tables 2-5 make it possible to conclude that DNSSEC validation is down, but before doing so we must examine whether there are other factors at work. Traffic in the November trace contained about 59 million questions,

while the January trace only had 37 million. In addition, the January trace was collected on a Sunday but the November one occurred on a Tuesday. In Table 6 we present more statistics about the two collection periods.

TABLE VI
STATISTICS FOR BOTH PERIODS

	Nov. 2010	Jan. 2011	% prior sample
Questions	58891887	37342902	62%
Resolvers	676599	573773	85%
No DNSSEC	214036	185954	86.8%
DNSSEC-Capable	462563	387779	83.8%
% DNSSEC-Capable	68.4	67.7	
Confirmed Validators	5621	4728	84%
Suspected Validators	106	785	

In both periods the percentage of confirmed validators is about 1.2 percent of the total number of resolvers, but that does not tell the whole story; for example, the lower query volume probably means we have less of a chance to see DS or DNSKEY queries from busy resolvers.

Looking back at dlv.isc.org we see over 7300 different sources doing look-ups, and less than half are on the list we built using our criteria.

D. How much validating goes on?

A more fundamental question is how many answers from .org are potentially validated as a fraction of all answers, which is partly answered by these facts:

- In the November 2010 sample, confirmed and possible validating resolvers accounted for 8 percent of the observed DNS queries to .org.
- In the January 2011 sample, confirmed and possible validating resolvers accounted for over 10 percent of observed DNS queries to .org.

Some of this traffic is caused by validators that advertise too large of a EDNS0 buffer yet cannot receive a DNSKEY RRset from .org that is larger than 1300 bytes. This is still an impressive percentage this early in DNSSEC deployment. However, note that some of this traffic may also be attributable to DDoS attacks that used the .org DNSKEY RRset as amplifier.

V. CONCLUSIONS

In this paper we analyzed traces from .org and tried to estimate the size of the DNSSEC validating resolver population, and had four proposed methods for identifying the validators. Due to the behavior of DNS resolvers we ruled out one of these techniques, since it will not work for validators that are "discovering" a domain or for busy resolvers that are "locked" in on a "close" authoritative server. We applied the other three techniques to the samples and were able to estimate the size of the validator population those samples represent. Based on our results it seems that by simply looking at DS queries it is possible to determine whether a resolver is a validator. This simplifies the search for validating resolvers as there is no need to correlate a DNSKEY query with a prior query. Ignoring DNSKEY has the further advantage that it eliminates interference from DDoS attacks that use DNSKEY RRsets.

The percentage of validation (8-10%) was higher than we expected; we think this is good news as it shows that validation is taking place and there are minimal reports of resolution problems due to validation taking too long. This shows that DNSSEC validation is a feasibilty for anyone.

REFERENCES

[1] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "DNS Security Introduction and Requirements." RFC 4033 (Proposed Standard), Mar. 2005. Updated by RFC 6014.

[2] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "Resource Records for the DNS Security Extensions." RFC 4034 (Proposed Standard), Mar. 2005. Updated by RFCs 4470, 6014.

[3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, "Protocol Modifications for the DNS Security Extensions." RFC 4035 (Proposed Standard), Mar. 2005. Updated by RFCs 4470, 6014.

[4] E. Osterweil, D. Massey, and L. Zhang, "Deploying and Monitoring DNS Security (DNSSEC)," *Computer Security Applications Conference, Annual*, pp. 429–438, 2009.

[5] J. Jung, E. Sit, H. Balakrishnan, and R. Morris, "DNS performance and the effectiveness of caching," in *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, IMW '01, (New York, NY, USA), pp. 153–167, ACM, 2001.

[6] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig, "Comparing DNS resolvers in the wild," in *Proceedings of the 10th annual conference on Internet measurement*, IMC '10, (New York, NY, USA), pp. 15–21, ACM, 2010.

[7] M. Larson and D. Blacka, "Port and Message ID Analysis of Resolvers Querying .com/.net Name Servers." OARC Workshop Fall 2008, Sep 2008. https://www.dns-oarc.net/files/workshop-2008/larson.pdf.

[8] G. Sisson, "Durz analysis." DNS-OARC Fall 2010 workshop, Mar 2010. http://bit.ly/dP1vGD.