

DNSSEC and Its Potential for DDoS Attacks

A Comprehensive Measurement Study

Roland van Rijswijk-Deij
University of Twente and
SURFnet bv
r.m.vanrijswijk@utwente.nl

Anna Sperotto
University of Twente
a.sperotto@utwente.nl

Aiko Pras
University of Twente
a.pras@utwente.nl

ABSTRACT

Over the past five years we have witnessed the introduction of DNSSEC, a security extension to the DNS that relies on digital signatures. DNSSEC strengthens DNS by preventing attacks such as cache poisoning. However, a common argument against the deployment of DNSSEC is its potential for abuse in Distributed Denial of Service (DDoS) attacks, in particular reflection and amplification attacks. DNS responses for a DNSSEC-signed domain are typically larger than those for an unsigned domain, thus, it may seem that DNSSEC could actually worsen the problem of DNS-based DDoS attacks. The potential for abuse in DNSSEC-signed domains has, however, never been assessed on a large scale.

In this paper we establish ground truth around this open question. We perform a detailed measurement on a large dataset of DNSSEC-signed domains, covering 70% (2.5 million) of all signed domains in operation today, and compare the potential for amplification attacks to a representative sample of domains without DNSSEC. At first glance, the outcome of these measurements confirms that DNSSEC indeed worsens the DDoS phenomenon. Closer examination, however, gives a more nuanced picture. DNSSEC really only makes the situation worse for one particular query type (ANY), for which responses may be over 50 times larger than the original query (and in rare cases up to 179×). We also discuss a number of mitigation strategies that can have immediate impact for operators and suggest future research directions with regards to these mitigation strategies.

Categories and Subject Descriptors

C.2 [Computer-Communication Networks]: Miscellaneous; C.4 [Performance of Systems]: Measurement Techniques

Keywords

DNS; DNSSEC; DDoS; amplification attack; reflection attack; measurements; denial-of-service; attack

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
IMC'14, November 5–7, 2014, Vancouver, BC, Canada.
Copyright 2014 ACM 978-1-4503-3213-2/14/11 ...\$15.00.
<http://dx.doi.org/10.1145/2663716.2663731>.

1. INTRODUCTION

Since they were first seen at scale at the turn of the century [1], Distributed Denial of Service (DDoS) attacks have become one of the biggest threats to the Internet's security and stability. The scale of DDoS attacks keeps growing; the "biggest DDoS ever" of 300 Gbit/s in early 2013¹ has already been surpassed this year by an attack that was 25% larger in volume².

These large volume attacks usually rely on the same basic principles. First, they use *spoofing*. Spoofing allows the attacker to falsify the source IP in a request to some network service, resulting in the response to this request being sent to the falsified source IP (this is known as *reflection*). Second, attackers leverage *amplification*, the principle that some network protocols return a large answer to a relatively small request. Amplification is of particular interest to an attacker since a small investment in attack traffic results in large attack volumes.

A commonly used DDoS attack is *DNS amplification* (as e.g. Arbor Networks' annual security report [2] shows). As the name suggests, this attack relies on bandwidth amplification using the DNS protocol, where amplification is defined as $\frac{\text{response size}}{\text{query size}}$. Typical DNS requests are in the order of magnitude of 20–60 bytes in size. The classic DNS protocol [3] limits responses to at most 512 bytes; assuming a request size of 40 bytes this already yields an amplification factor of $\frac{512}{40} \approx 12.8$. More recent extensions to DNS that allow for larger responses easily result in amplification factors of 100 or more.

Over the past six years, since Dan Kaminsky disclosed a serious vulnerability in the DNS protocol [4], a major overhaul of the DNS has been underway: the introduction of the DNS Security Extensions (DNSSEC) protocol. This deployment is now bearing fruit; where in 2008 there were a handful of DNSSEC-signed domains, the total number now tops 3.5 million³. One of the key features of the DNSSEC protocol is the introduction of digital signatures in DNS responses. Consequently, the size of DNS responses increases. Many experts consider this a major drawback of DNSSEC (e.g. Cowperthwaite & Somayaji [5]) and noted opponents of DNSSEC cite this as one of the reasons why they feel DNSSEC should not be used (e.g. Bernstein [6]).

Although the DDoS potential keeps coming up time and again in discussions about DNSSEC, very little ground truth

¹<http://www.bbc.com/news/technology-21954636>

²<http://www.bbc.com/news/technology-26136774>

³based on statistical sources listed at <http://www.internetsociety.org/deploy360/dnssec/statistics/>

exists about the actual DDoS potential of the millions of DNSSEC-signed domains that are online today. This leads us to the main question we will answer in this paper: *How bad is DNSSEC really?*

Contribution.

In this paper, we provide the first comprehensive measurement of the DDoS potential of DNSSEC. Our measurements encompass 70% of all DNSSEC-signed domains in operation today and is based on data from six major top-level domains (TLDs) each with a significant number (> 10K) of signed domains. We compare the DDoS potential of DNSSEC-signed domains to a representative sample of domains without DNSSEC. Analysis of our measurements shows that the average amplification of DNSSEC exceeds that of regular DNS many times (by a factor of 6x - 12x). At first glance, this is worrying. Looking deeper, however, it becomes evident that extreme amplification only occurs for a certain type of query (ANY) that is often abused for amplification attacks. For “normal” DNS queries, the picture is much more nuanced and does not warrant some claims that DNSSEC should not be deployed [5, 6]. Next to that, a number of measures can be taken to dampen the DDoS potential of DNSSEC significantly, although further work in this area is required.

To encourage further study we make the data collected for our measurements available as open data to the Internet measurement community (see Sec. 8).

Organization of this paper.

This paper is organized as follows: Sec. 2 provides background material on DNS amplification as an attack method. Sec. 3 describes our methodology and measurement setup. In Sec. 4 we discuss the data sets collected during our measurements. Sec. 5 gives a detailed analysis of the data we obtained. Sec. 6 provides an overview of current countermeasures based on the literature and calls for the introduction of additional countermeasures. We discuss related work in Sec. 7. Finally, we draw conclusions and suggest future research directions in Sec. 8.

2. BACKGROUND

2.1 DNS amplification: a brief primer

2.1.1 Attacks using open resolvers

Fig. 1 shows how DNS amplification attacks work. Attacks are initiated from a swarm of machines (left-hand side of the figure) under the control of the attacker. The attacker uses this swarm to send large numbers of DNS queries in which the sender IP address is spoofed to be the victim’s IP address (bottom middle of the figure). Queries are sent ① to so-called *open DNS resolvers*. These are misconfigured DNS resolvers that do not restrict which clients are allowed to send them queries. In turn, the open resolvers will – if the query result is not cached – contact the appropriate authoritative DNS servers ② to resolve the query. Finally, the open resolvers will send the responses ③ to the victim. In general, the queries ① are small whereas the responses ③ are large, hence achieving amplification.

Unfortunately, open resolvers are plentiful on the Internet. Kührer et al. [7] report observing between 23 and 25.5 million open resolvers during weekly Internet-wide scans over a

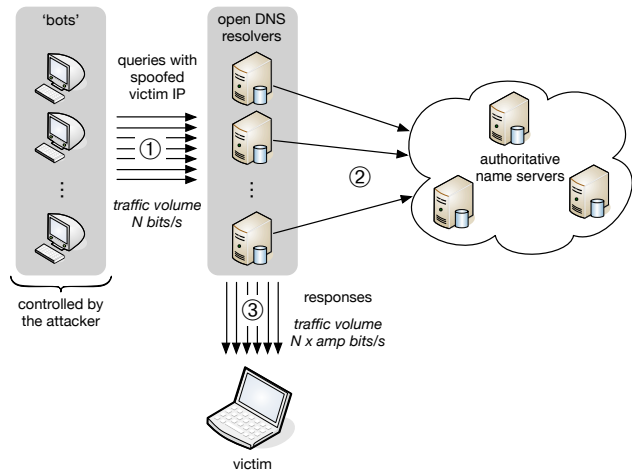


Figure 1: DNS amplification attack

period of 4 months between November 2013 and February 2014. This makes this attack easy to carry out and thus attractive for attackers.

2.1.2 DNS extensions (EDNS0)

Typically, the goal of an attacker is to achieve a high amplification factor. This gives him the best return on investment where he only needs to generate a small amount of traffic for a large attack. With the introduction of the EDNS0 extension [8] larger DNS responses (than the original 512 bytes) become possible. EDNS0 allows clients and servers to specify the maximum response size they support. Fig. 2 shows whether clients to one of SURFnet’s⁴ authoritative name servers use EDNS0. As the figure shows, about two thirds of clients use EDNS0; earlier research (e.g. [9]) shows that this figure is similar for other name servers across the Internet. The prevalent configuration for EDNS0 is to set the maximum response size to approximately 4Kbytes. Fig. 3 shows the maximum response size reported by clients of the same name server as in Fig. 2, with around 90% advertising a size over 4000 bytes. Again, this matches earlier results reported by Kreibich et al. [9].

⁴The National Research and Education Network in the Netherlands, <http://www.surf.nl/en/about-surf/subsidiaries/surfnet>

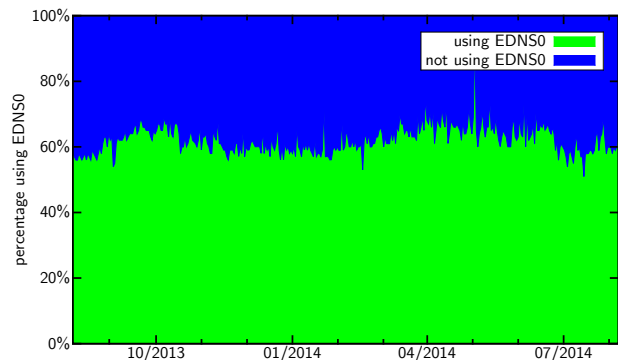


Figure 2: EDNS0 use for ns1.surfnet.nl

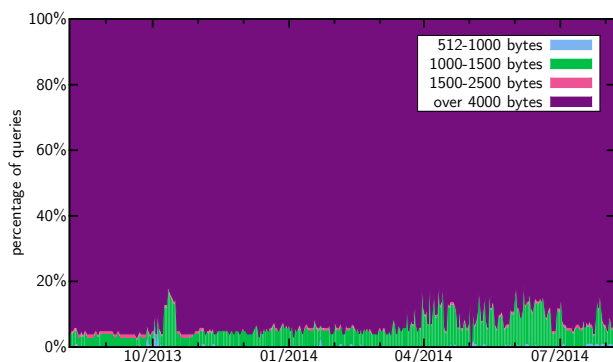


Figure 3: Advertised EDNS0 maximum response size for ns1.surfnet.nl

In theory, given a query size of 40 bytes, an amplification of $\frac{4096}{40} \approx 102.4$ is achievable. Table 1 gives some examples of the theoretical attack volumes that can be achieved if the attacker generates 100Mbit/s in queries for various amplification factors.

Q-size (bytes)	R-size (bytes)	Ampl. (factor)	Attacker (bits/s)	Victim (bits/s)
40	512	12.8	100M	1.28G
40	1472	36.8	100M	3.68G
40	4096	102.4	100M	10.24G

Table 1: Theoretical effect of DNS amplification

As the table shows, when using EDNS0’s expanded maximum response size, significant attack results can be achieved with a relatively small investment in bandwidth on the side of the attacker.

2.1.3 Crafted domains

Also of interest is how attackers select the DNS query to use in the attack. One current practice is that attackers craft special domains for which certain DNS queries are guaranteed to return large responses. Vaughn & Evron [10], for instance, analyse a case where an attacker crafted a domain with a large TXT record that guaranteed a large response to certain queries. A particularly interesting query type for attackers to abuse is the ANY query; when faced with such a query, name servers are supposed to return all record types for the name specified in the query. In practice, a combination of both a crafted domain as well as the use of ANY queries are often observed. In our own research group, as part of another project in which we are investigating “DDoS-for-hire” services [11] we have observed a number of DNS amplification attacks, most of which used EDNS0, ANY queries and crafted domains.

2.2 Benefits DNSSEC offers to the attacker

There are a number of “hinderances” for attackers given the attack methodology described in the previous subsections:

- bandwidth on the open resolvers that are abused for the attack may be scarce (for instance because the open resolver is on a DSL CPE device), necessitating the use of large numbers of amplifiers;

- creating crafted domains makes the attacker vulnerable to prosecution since domain registrations can – sometimes – be traced back to the attacker;
- crafted domains may be taken down when discovered;
- filtering of responses for crafted domains can be a possible mitigation strategy for victims.

DNSSEC-signed domains offer an attractive alternative to attackers since they help avoid the hinderances described above. Given the high number of signed domains, it gets more interesting to directly abuse authoritative name servers for DNSSEC-signed domains, instead of open resolvers; it is more likely that authoritative name servers have sufficient bandwidth available than e.g. open resolvers on DSL CPE devices. Also, instead of crafting a dedicated domain, attackers can now choose from a wide range of DNSSEC-signed domains for which the answers to e.g. ANY queries are most likely quite large due to the inclusion of DNSSEC key material and signatures. This makes attackers less vulnerable to discovery, prosecution and take-down of the crafted domains. Finally, DNS responses for legitimate DNSSEC-signed domains are much harder to filter since filtering is likely to also impact legitimate queries and responses.

These benefits to attackers underline the necessity of establishing ground truth about the actual DDoS potential of the already large and growing number of DNSSEC-signed domains. With this data in hand, it becomes possible to develop targeted strategies to address this potentially dangerous side effect of deploying DNSSEC.

The importance of studying this topic now is also emphasised by the fact that there have already been attacks that abuse DNSSEC-signed domains for amplification. In mid-2012 many organizations across the Internet saw their DNSSEC-signed domains being abused directly on authoritative name servers for amplification attacks. Documentation about these attacks is, unfortunately, not available as they were mainly discussed in private by large operators, on mailing lists and at meetings such as RIPE and the IETF. The lead author of this paper, however, was involved in many of these discussions. An indirect discussion of these attacks was written up by Cisco in their security blog⁵.

We note that the DNS and DNSSEC community (consisting of network engineers, DNS software developers and organisations concerned with online security) has already started to take steps to mitigate the dangers of DNSSEC-based amplification attacks by introducing Response Rate Limiting (RRL), which is discussed in Sec. 6.

3. METHODOLOGY

3.1 High-level goals

Our high-level goal, as stated in Sec. 1, is to gauge the DDoS potential of DNSSEC-signed domains in comparison to unsigned domains. In order to achieve this goal, we perform a large scale measurement of all DNSSEC-signed domains in six top-level domains (TLDs) that have large numbers of signed domains. We compare these against a representative sample of unsigned domains in the same TLDs.

⁵Case study 2 in <http://blogs.cisco.com/security/real-world-dns-abuse-finding-common-ground/>

3.2 Acceptable upper limit for amplification

An important question to ask – and one that is difficult to answer – is what is an acceptable upper limit on the amplification possible for a particular domain on a particular server. The simple answer is that any amplification is bad, but such an answer is too simple, since the DNS protocol already inherently has some amplification effect.

As we outlined in the previous section, EDNS0 vastly increases the amplification potential in DNS. Since – other than combating open resolvers – no large scale action has been taken to reduce the amplification potential in pre-EDNS0 classic DNS, we take the maximum amplification achievable using classic DNS as an upper limit. We calculate⁶ this maximum achievable amplification by 1) assuming a query for the shortest name in a domain (e.g. “x.com”), which will consequently result in the smallest query and 2) a response that uses the maximum allowed size for regular DNS (512 bytes). For our example (x.com) this yields an amplification of $\frac{\text{response size}}{\text{query size}} = \frac{512}{23} \approx 22.3$.

3.3 Query types

We selected a number of queries to perform for each domain. These queries are specified below, together with a rationale for why we selected this particular query type:

- **ANY** – as explained in Sec. 2.1.2 this query results in the largest possible response since it includes all resource records for the queried name.
- **MX** – returns the names of mail exchangers for the domain; domains usually have more than one mail exchanger hence the answer to this query may be relatively large.
- **NS** – returns all authoritative name servers for a domain; again, domains usually have more than one authoritative name server, thus the answer to this query may also be relatively large.
- **A** – returns the IPv4 address(es) for the queried name and is typically the most common DNS query performed for a domain.
- **AAAA** – returns the IPv6 address(es) for the queried name and is another common query type since most modern software will look for both the IPv4 and the IPv6 address of a host.
- **TXT** – returns textual information for the queried name; this can, for instance, be information about mail handling for the domain (used by spam filters). Additionally, as mentioned in Sec. 2.1.3 this query type is sometimes used by attackers in domains especially crafted for amplification attacks.

For the latter three query types (**A**, **AAAA** and **TXT**) we perform multiple queries. One for the so-called apex record (denoted by **@**), one for the name **www** under the domain and one for the name **mail**. We expect that at least one of these names exists in most domains.

⁶Note that this is not just a theoretical value; analysis of our measurements shows that each measured TLD has a non-negligible number of domains that meet this pattern (small query, maximum size response) when queried using classic DNS.

For DNSSEC-signed domains we also measure the answer to two DNSSEC-specific query types:

- **DNSKEY** – returns the set of public keys required to validate signatures in a domain. The answer is usually large as there is often more than one key and individual key records are large as RSA keys of sizes 1024- and 2048-bits are commonly used (as suggested in [12] and because these are default values for many DNSSEC implementations).
- **NSEC(3)** – this is not a query type as such, but is the record type that is returned when the queried name does not exist; **NSEC(3)** is also known as *authenticated denial of existence* and proves with a digital signature that the queried name does not exist. Especially responses of the newer variant **NSEC3** are likely to be quite large.

Each individual query is performed once using classic DNS and once using EDNS0 with the EDNS0 maximum response size set to 32768⁷. For DNSSEC-signed domains we also perform the query with EDNS0 and the DNSSEC OK (DO) flag set to **true**, to get DNSSEC-signed responses.

3.4 Metrics

The metrics we record for each query are:

- the query size (DNS UDP datagram size);
- the response size (DNS UDP datagram size, possibly reassembled from multiple UDP fragments);
- the amplification factor (defined as $\frac{\text{response size}}{\text{query size}}$);
- the EDNS0 maximum response size reported by the authoritative name server (puts an upper bound on the maximum amplification that can be achieved using this particular name server);
- whether or not the response was truncated (this indicates that the authoritative name server was unable or unwilling to return all requested data in the response), and indicates that the querying host should fall back to TCP to get the full response;
- the number of answers in the response;
- the number of authority records in the response (in most responses including this information is optional and indicates which name servers are authoritative for the queried name);
- the number of additional records in the response (these optional records can for instance specify IP addresses for name servers included in the authority section);
- the number of distinct resource record types in the response.

⁷We chose this value to also register results that exceed the commonly used maximum response size of 4KB; we decided not to use the maximum value (65535) since we did not want to risk running into possible boundary conditions in DNS software implementations.

To ensure that we are not inadvertently including domains of one kind in the other data set, we check for each domain whether it is DNSSEC-signed or not (by checking the presence of `DNSKEY` records). If the domain turns out to belong to another category than expected, we exclude this domain from the measurement. Equally, those domains for which we are unable to determine the set of authoritative name servers or for which none of the authoritative name servers responds to queries are excluded from the data sets. We note that, for reasons of efficiency, we attempt each query only once, since authoritative name servers are generally always online.

3.5 Measurement software

To perform the measurements we developed two applications. The first is a zone file parser, which given the size of the DNS zones for TLDs, we needed specially tailored software for. The parser application is able to extract DNSSEC secure delegations, i.e. DNS delegations for which a DNSSEC chain of trust is established by including one or more `DS` records in the TLD (see Sec. 5 of [13]). It is also able to extract regular domains from the zone. The parser stores the extracted domain names in a SQLite database. Optionally, the application can take a random sample of domains.

The second application is a scanner application. This application operates on the database created by the zone file parser and for each domain in the database will perform the queries outlined in Sec. 3.4. The application is massively parallel and launches several hundreds of threads to perform the scan. Scanning takes place in two phases. In phase one, the application determines the set of authoritative name servers for each domain in the database and the corresponding IPv4 and IPv6 addresses for these name servers. In phase two, the application sends the queries described in Sec. 3.4 to each individual IP address for each name server of each domain. After each set of queries to a single IP address, the domain is placed back at the end of the queue until no more IP addresses remain for the domain to be scanned in which case it is marked as completed. This design choice ensures that no excessive amounts of queries for a single domain are sent in bursts while keeping the amount of state the scanner needs to maintain manageable. We found that this strategy strikes an optimal balance between the load imposed on the scanning system as well as the scanned systems.

Both applications make extensive use of the LDNS library by NLnet Labs⁸, a popular software package that provides a comprehensive set of DNS-related functions.

Note that the SQLite database schema was designed such that the data can easily be anonymized. This was done on purpose because the resulting data sets will be shared as open data with the research community.

3.6 Ethical considerations

As mentioned in the previous section, we took particular care to ensure that our measurements do not impose an undue burden on authoritative name servers we scan. To this end we made deliberate design choices as discussed in Sec. 3.5. In addition to this, we monitored our experiment while it was running to check if the design worked as foreseen. Additionally, we only performed legitimate queries

that are expected to be part of day-to-day traffic to authoritative name servers.

With respect to the data sets for the TLDs we used, we obtained these through specific processes established by the TLD operator or under a specific contract with individual TLDs. In both cases, we took care to inform the TLDs about the purpose for which we were going to use the data and obtained their consent.

4. DATA SETS

4.1 Source data

TLD	Data obtained	#domains	#DNSSEC	
.com	Full zone	113.1M	326.5k	(0.3%)
.net	Full zone	15.2M	69.5k	(0.5%)
.org	Full zone	10.3M	37.6k	(0.4%)
.nl	Selection	5.4M	1696.1k	(31.2%)
.se	Full zone	1.4M	334.9k	(24.8%)
.uk	Selection	10.6M	10.2k	(0.1%)

Table 2: Overview of source data

We obtained data covering 70% of all 3.5 million DNSSEC-signed domains from six different top-level domains. Tab. 2 lists the TLDs from which we obtained data. The table lists the type of data obtained (either the full zone or a selection containing all secure delegations and a random sample of unsigned domains), the total number of domains in the TLD and the number of secure delegations (as an absolute value and a percentage).

4.2 Collected data

TLD	#domains	#failed	#skipped	#queried	#queries	#auth ns
.com	326504	7416	471	318576	54.6 M	2550
.net	69552	2672	55	66814	11.0 M	2476
.org	37621	555	19	37024	6.7 M	2073
.nl	1696103	12304	1002	1682770	233.3 M	1316
.se	334880	8696	100	326067	43.3 M	3681
.uk	10225	314	10	9894	1.6 M	570

Table 3: Overview of DNSSEC data sets

TLD	#domains	#failed	#skipped	#queried	#queries	#auth ns
.com	498502	55909	2231	436593	37.6 M	72168
.net	99564	13904	355	84882	7.4 M	26396
.org	100000	11031	277	88372	7.5 M	27761
.nl	1000000	69092	6812	921441	69.3 M	31108
.se	499999	37361	149560	311871	21.5 M	23756
.uk	26131	3883	92	21858	1.6 M	7091

Table 4: Overview of non-DNSSEC data sets

We ran two scans for each TLD in the source data set. The first scan covered all DNSSEC-signed domains in the TLD. The second scan examined a representative uniformly random sample of unsigned domains with a size in the same order of magnitude as the number of DNSSEC-signed domains in the TLD.

For each scan type Tab. 3 and Tab. 4 show the total number of domains for which queries were attempted, the number of domains for which we failed to obtain the list of authoritative name servers, the number of domains that were skipped (because they were DNSSEC-signed whereas they

⁸<http://www.nlnetlabs.nl/projects/ldns/>

were expected not to be or vice versa), the actual number of domains that were successfully queried, the total number of queries included in the data set and the number of distinct authoritative name servers observed during the scan. We note that there may be a slight difference between the total number of domains for which queries were attempted (col. 2) and the number of failed, skipped and successfully queried domains added up (col. 3 + 4 + 5). This is because for a small number of domains although we were able to determine the set of authoritative name servers none of these responded to queries. Since for both the regular as well as the DNSSEC data sets this difference is very small (0.52% and 0.03% on average over all TLDs respectively) it is not shown in the table.

4.3 General observations

If we look at the result data sets in Tab. 3 and Tab. 4 three observations stand out.

1. The fraction of domains for which the set of authoritative name servers could not be determined is significantly larger for unsigned domains (on average 10.9% for unsigned domains versus 2.3% for DNSSEC-signed domains). This seems to imply that DNSSEC-signed domains are generally configured better and are more likely to have functioning delegations.
2. The ratio of distinct authoritative name servers versus the number of domains in the data set is much higher for unsigned domains (by more than a factor of 10). This is probably due to the fact that a number of large DNS operators enabled DNSSEC for all domains they manage stimulated by TLDs like .nl, .se and .org offering financial incentives for deploying DNSSEC.
3. A very high number of domains in the .se TLD were skipped in the non-DNSSEC data set because they turned out to be DNSSEC-signed after all; a random sampling of the skipped domains shows that this is very likely due to one large DNS operator having enabled DNSSEC without creating secure delegations in the parent TLD.

5. ANALYSIS

5.1 Introduction

In this section, we provide a detailed analysis of our measurements and examine the differences between regular and DNSSEC-signed domains. Rather than looking at the exact response size in bytes for different queries, we use the amplification factor as main metric. We do this because it allows for a non-biased comparison between queries for different domains independent of the length of the domain name. To illustrate this with an example, the shortest scanned domain name in .com is 1 character long; the query size for this domain is 34 bytes, the response size for an ANY query with DNSSEC enabled is 3549 bytes. This gives an amplification factor of 104.4, which is very high as we will see later. The query size for the longest name (63 characters) scanned in .com is 96 bytes and the response size is 3805 bytes. That only yields an amplification factor of 39.6 (around average as we will see later). If we had compared these two domains on the basis of the response size then the second domain would have been classified as “worse” than the first domain,

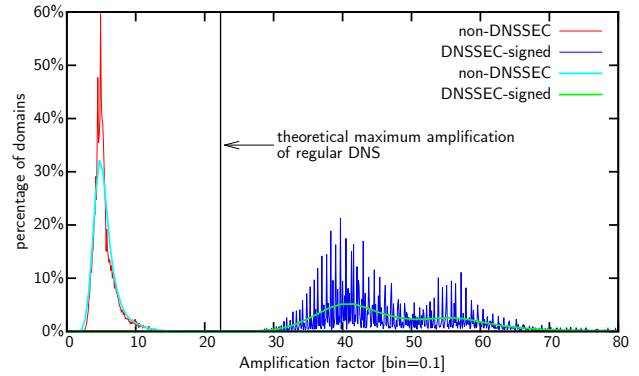


Figure 4: Amplification of ANY queries for .net

whereas in fact it is not. We also note that the amplification factor more accurately reflects the return on investment for attackers. A certain amount of traffic n sent by an attacker results in an attack that is $n \times \text{amplification}$ in size. Note that we calculate the amplification factor using the UDP datagram size for both query and response rather than using the full packet size. This prevents noise from varying packet sizes caused by differences in the underlying network technology.

Another general remark concerns the selection criteria for the graphs we plot. In most of our analyses, we only include responses we consider valid, defined as those responses that include at least one answer in the answer section and that have a response code (RCODE) that indicates success (NOERROR). We do this to exclude as many authenticated denial-of-existence answers as possible from the data sets used to plot the graphs since these may skew the data in the graphs (as especially the A and AAAA queries we perform may result in authenticated denial-of-existence as the name we query for may not exist). We cover authenticated denial-of-existence itself in a separate graph, which of course does include responses with no answers and with response codes that indicate that the name does not exist.

In our analysis, we look at three areas of interest. First we examine the main vehicle for DNS amplification attacks, ANY queries. Next, we look at query types for which we expect large answers. We finish by examining the bread and butter of the DNS, address queries.

5.2 ANY amplification increase

Since ANY queries are the most interesting for attackers, we start our analysis by looking at these in more detail. To compare the amplification factor that is achievable with an ANY query for regular domains and DNSSEC-signed domains we plot this factor for both types of domains in a single graph. In order to do this, we distribute the data across bins of size 0.1; plotting the actual data results in a graph as shown in Fig. 4, which shows the ANY amplification factor for the .net TLD. As can be seen, graphing the original data (thin red and blue lines) leads to a jagged graph. This is due to the fact that the amplification factor is derived from discrete values (query and response size) resulting in a bias for certain bins. To mitigate this, we have chosen to represent the data using a Bézier curve that follows the average fill of the bins (thicker cyan and green lines).

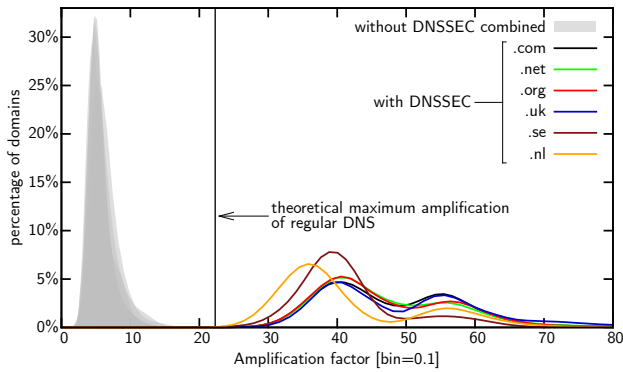


Figure 5: Amplification of ANY queries (all TLDs)

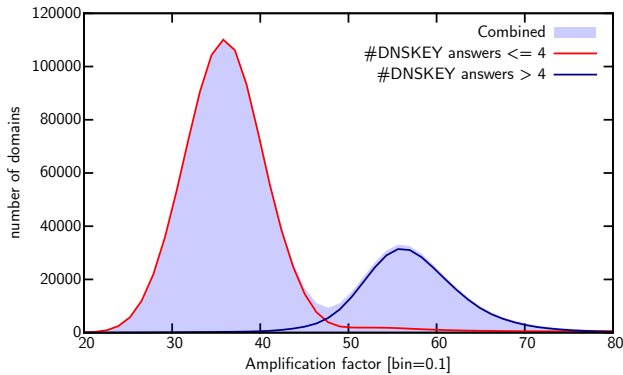


Figure 6: Examining the two peaks in .nl

The graph in Fig. 5, plotted as described above, compares the ANY amplification factor for all TLDs for which we performed measurements. Regular domains are represented by filled grey curves whereas DNSSEC-signed domains are represented using coloured lines. This approach to plotting is repeated in all other graphs that compare amplification of regular versus DNSSEC-signed domains. The vertical black line (indicated with an arrow) represents the theoretical maximum achievable amplification factor using regular DNS, which we set as an acceptable upper limit to amplification in Sec. 3.2.

Looking at Fig. 5 we see that, as expected, the amplification factor for DNSSEC-signed domains is significantly higher than for unsigned domains. On average, the amplification for unsigned domains is around 5.9 whereas for DNSSEC-signed domains the average lies around 47.2 (about $8\times$ higher). We also note that – apart from outliers not visible in the graph – the amplification factor for ANY queries for DNSSEC-signed domains always exceeds the acceptable upper limit we set.

Next, the distribution of the amplification factors observed is much more spread out for DNSSEC-signed domains. This is as expected; a DNSSEC-signed answer includes a separate signature (RRSIG record) for each resource record set (all records of a certain type for a name). Thus, the more different resource record types in an answer the more signatures. Signature records are large (> 150 bytes for a 1024-bit signature) and contribute significantly to the amplification factor.

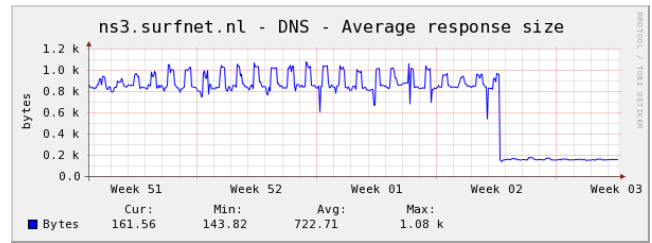


Figure 7: Effect of disabling optional response sections on ns3.surfnet.nl

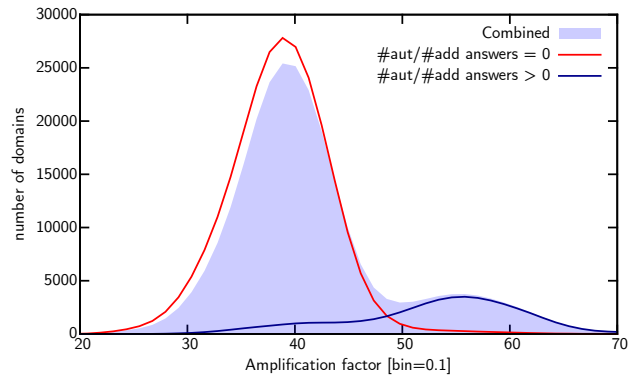


Figure 8: Examining the two peaks in .se

Finally we note that the distribution of amplification factors for DNSSEC-signed domains shows two peaks for each TLD, one around 40 and a lower peak around 55. We found two effects that play a role in this.

The first effect occurs for the majority of the TLDs we measured (.com, .net, .org, .nl and .uk). For these TLDs, the two peaks can be explained by looking at the number of DNSKEY records in the zone. The first peak coincides with having 2 DNSKEY records in the zone, the second peak coincides with having 3 or more DNSKEY records in the zone. Fig. 6 shows this correlation for the .nl TLD (the blue line in Fig. 6 corresponds to the data for .nl as graphed in Fig. 5 in orange). Note that we do not have direct data for the number of DNSKEY records per scanned domain. Rather, we derive this from the number of answers in the DNSKEY query for a domain. If this number is 4 or below (red line in Fig. 6), then there are 2 or less DNSKEYs present (4 answers means 2 DNSKEY records and 2 signatures in the answer, one signature with each key). If this number is higher than 4 (blue line in Fig. 6) then there are 3 or more DNSKEYs present. The reason for having different numbers of DNSKEY records in a domain is simple: different software implementations that use different key rollover strategies (a discussion of which is outside of the scope of this paper, for more information on DNSSEC key rollovers see e.g. [14]).

The second effect occurs for one TLD (.se) only. Interestingly, in this case the two peaks show no relation to the number of DNSKEY records. Rather, the peaks are related to whether or not the additional and authority sections in a DNS answer are filled. These sections are optional in most DNS responses (see Sec. 4 of [3]). Most DNS software implementations have a configuration option that allows administrators to disable filling of these optional sections of a DNS

answer⁹. Disabling these optional responses can have a dramatic effect on the response size. Fig. 7 shows the change in average response size (a reduction of $\pm 80\%$) for an authoritative name server. Especially for DNSSEC-signed domains, the difference in size can be dramatic since signatures will be included in these optional sections. In Fig. 8 we graphed the amplification for DNS responses that did not include the optional authority and additional sections (red line), the amplification for responses that did include optional sections (blue line) and the original data for ANY responses in .se (blue filled curve) corresponding to the brown line in Fig. 5.

5.2.1 Outliers

Fig. 5 shows the distribution in amplification factor for the majority of domains measured. There are, however, outliers that have much higher amplification factors. Tab. 5 gives an overview of the outliers. Per TLD, it shows the number of domains with an amplification factor higher than 100 as well as the absolute maximum amplification factor of any domain measured.

TLD	#amp. > 100	%	maximum amplification	
			with DNSSEC	w/o DNSSEC
.com	144	(0.05%)	119.2	75.0
.net	168	(0.25%)	178.6	51.0
.org	139	(0.38%)	143.1	33.6
.nl	145	(0.01%)	131.0	80.9
.se	211	(0.06%)	120.0	63.7
.uk	26	(0.26%)	148.6	26.7

Table 5: Outliers per TLD

The table shows some staggering outliers, especially keeping in mind that an amplification factor of 100 means an attacker can mount an attack of 1Gbit/s by sending only 10Mbit/s. For comparison, the rightmost column of Tab. 5 shows the maximum amplification factor for the unsigned domains we measured. These are clearly a lot lower, with none exceeding 80.9.

5.3 Other large queries (DNSKEY, NSEC3, MX, NS, TXT)

First, we analyse two DNSSEC-specific query types, the DNSKEY query and authenticated denial-of-existence. As explained in Sec. 3.4, the answer to a DNSKEY query for a domain is likely to be large. Fig. 9 shows the amplification factor for DNSKEY queries for all TLDs we measured. On average 37.8% of DNSKEY queries yield an amplification factor that exceeds the acceptable upper limit we set in Sec. 3.2. We observe that most TLDs have two peaks in the graph; this is again caused by varying numbers of DNSKEY records per domain as described in the analysis of ANY queries in Sec. 5.2.

The other query type likely to result in large responses when DNSSEC is used, is a query that has no answer. Just like in regular DNS, a special response is returned when the name and/or record type queried for does not exist. The difference is that in DNSSEC a proof is included that the name and/or record type does not exist. This is called authenticated denial-of-existence and relies on two special record

⁹For instance, the popular DNS implementation BIND has the `minimal-responses` option, for more information see <ftp://ftp.isc.org/isc/bind9/cur/9.10/doc/arm/>

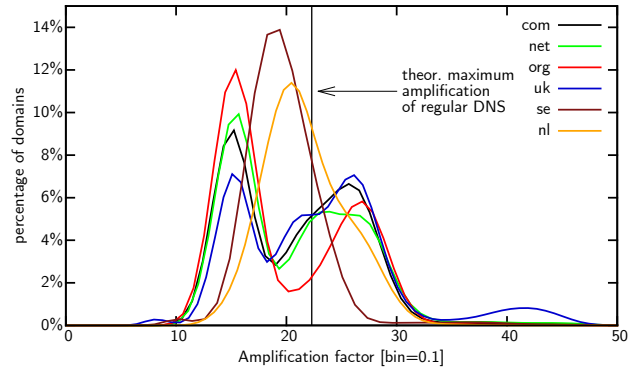


Figure 9: Amplification factor of DNSKEY queries

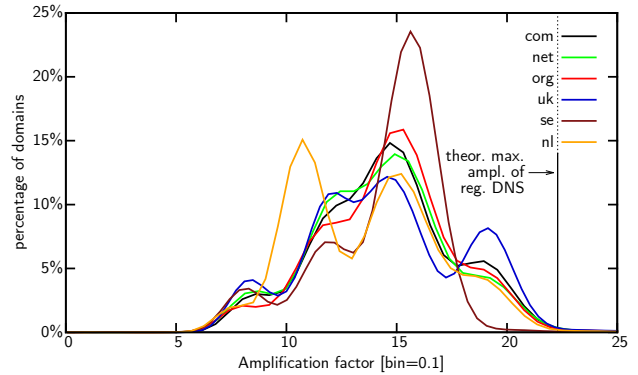


Figure 10: Amplification factor of authenticated denial-of-existence

types, NSEC and NSEC3. These record types are two protocol variants that achieve the same goal. The difference between the two is that NSEC uses names in proofs whereas NSEC3 uses hash values in proofs (a detailed discussion of the protocol differences is out of scope for this paper). Depending on the protocol and the type of proof required a number of NSEC or NSEC3 records are included in an authenticated denial-of-existence proof, each accompanied by a signature. Fig. 10 shows the amplification factor for NSEC(3) responses for all TLDs measured. As the graph shows, except for some outliers, the majority of responses fall within the acceptable upper limit we set. The graph contains multiple peaks for each TLD. Analysis of the data shows that the peaks are related to the number of NSEC(3) records in a response.

In Sec. 3.4 on metrics we said that we expected two other types of queries, the MX and NS query types, to also yield potentially large answers. Our measurements show, however, that this is not the case. Fig. 11 and Fig. 12 respectively show a comparison between the amplification factors we measured for regular domains versus DNSSEC-signed domains for MX and NS queries. As both graphs illustrate, the increase in amplification is low (between a 2 \times and 3 \times) and – outliers excepted – remains well below the acceptable maximum upper limit we set in Sec. 3.2. This is in line with other regular queries that will be discussed in the next section.

Similarly, we discussed that the TXT query type, that is for instance commonly used to convey information about spam filtering, is sometime abused by attackers in specially crafted

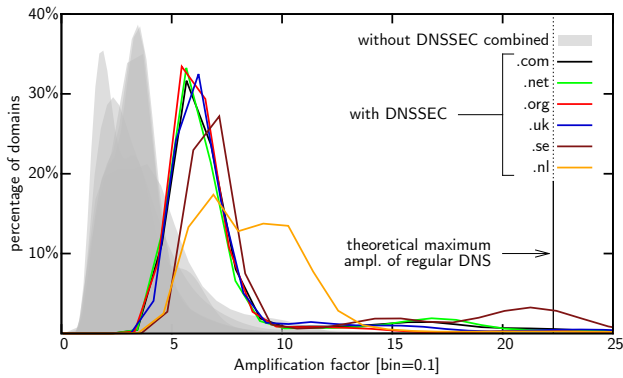


Figure 11: Amplification of MX queries

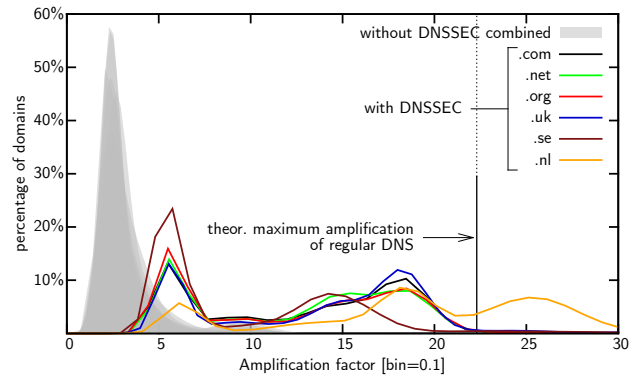


Figure 13: Amplification of TXT queries

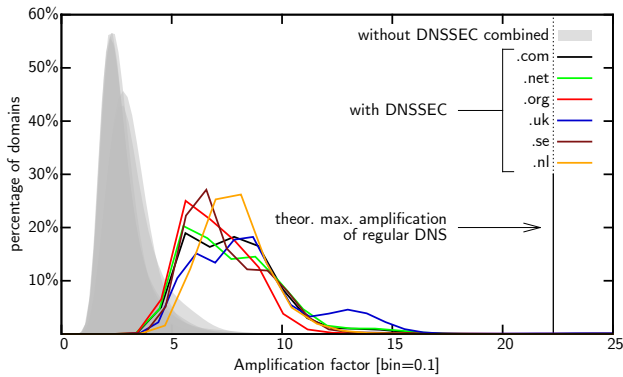


Figure 12: Amplification of NS queries

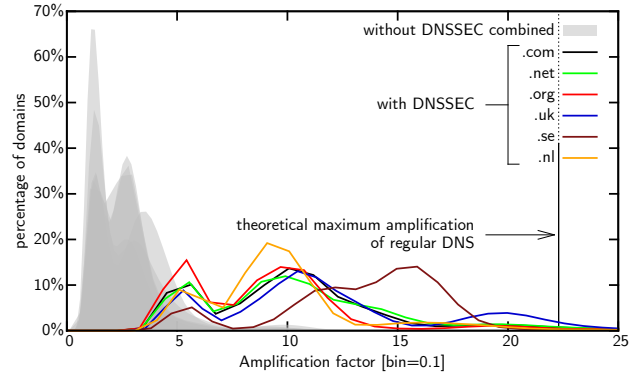


Figure 14: Amplification of A queries

domains created to perform DDoS attacks (see Sec. 2.1.3). Fig. 13 shows the comparison in amplification for TXT queries. As the graph shows, the majority of responses is below the acceptable upper limit. We note that for the .nl domain there is a small peak in responses just beyond the acceptable upper limit. Closer examination shows that this peak is caused by authenticated denial-of-existence answers. Recall from Sec. 5.1 that we only include responses in the data set that have one or more actual answers in them. It turns out that the responses that make up this peak are CNAME responses. In DNS, a CNAME functions like an alias that refers to another name. The DNS server will attempt to expand a CNAME answer to include the actual record that the CNAME refers to. If this is not possible (because no record of the queried type exists) then it will return only the CNAME answer together with an authenticated denial-of-existence (NSEC(3)) proof in the authority section of the response. We examined the data set used to plot Fig. 13 and compared the average number of answers in the authority section for responses both below as well as above the acceptable upper limit. For responses below the limit there are about 2.98 answers in the authority section whereas for responses above the upper limit this number is 7.75. Examination of a random sample of domains combined with this information provides a clear indication that this peak is caused by authenticated denial-of-existence answers related to CNAME expansion. We repeated this analysis for the other TLDs and saw a similar pattern albeit with peaks at slightly lower amplification factors (between 15 and 20).

5.4 Regular queries (A, AAAA)

We end with queries for the IPv4 (A) and IPv6 (AAAA) addresses of a name, which are the bread and butter of DNS. We start with the first, A queries. Fig. 14 shows the comparison in amplification factor between regular and DNSSEC-signed domains. As the graph shows, the vast majority of responses to A queries is well below the acceptable upper limit defined in Sec. 3.2. Generally speaking, the increase in amplification for DNSSEC-signed domains is between $2\times$ and $4\times$. We examined the higher amplification factors, above 12.5 present in the graph, to find out why these responses are larger. We believe that this is caused by a configuration difference in the name servers. Responses with an amplification factor above 12.5 include significantly more answers in the authority and additional section (3.90 and 3.46 for responses with amplification > 12.5 , 0.16 and 0.02 for responses with amplification < 12.5). This indicates that answers with an amplification > 12.5 come from name servers that fill the optional authority and additional sections, whereas answers with a lower amplification come from name servers configured to give minimal responses (see also Sec. 5.2).

Fig. 15 shows the situation for AAAA queries. The graph shows strong similarities with the graph for TXT records, with the majority of responses falling within the acceptable upper limit and with a small peak for the .nl domain just beyond the upper limit around an amplification factor of 26. Repeating the analysis we applied to TXT and A records shows that the cause of this peak is the same as for TXT records, namely authenticated denial-of-existence answers for CNAME

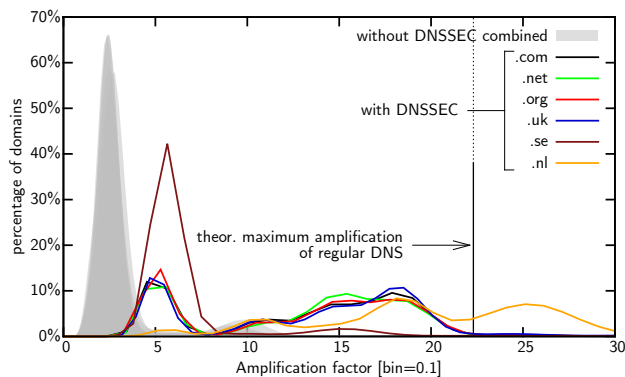


Figure 15: Amplification of AAAA queries

expansions. Again, this also proved to be the case for the other TLDs which have a similar peak around the somewhat lower amplification factor of 18.

6. COUNTERMEASURES

In Sec. 5 we showed that certain queries for DNSSEC-signed domains yield amplification factors that allow attackers to attain high attack volumes with a relatively small investment in bandwidth on their side. In this section we discuss a number of mitigation strategies to address this problem.

6.1 Ingress filtering

The root cause of any type of amplification attack is attackers having the ability to spoof source IP addresses in connectionless protocols (e.g. UDP). This problem has been known for a long time and the obvious solution is to stop source address spoofing. The solution is known as ingress filtering and is described in BCP 38 [15]. The basic premise behind ingress filtering is that network operators only allow traffic to enter their network if the source IP address is a legitimate address within their network. Although this solution is highly effective, it only works if deployed Internet-wide. Unfortunately, that is not the case. NPS runs a project that measures deployment of BCP 38¹⁰. Although their statistics indicate that the majority of networks implement ingress filtering, a large number still do not, making this mitigation strategy ineffective.

6.2 Response Rate Limiting

As mentioned briefly in Sec. 2.2, Response Rate Limiting (RRL) was introduced in 2012 after a slew of DNS amplification attacks that abused DNSSEC-signed domains. The idea behind RRL is that authoritative name servers rate limit outgoing responses, when responses to the same query are sent repeatedly to the same IP block¹¹. At first glance, RRL can significantly dampen the impact of DNS amplification attacks. A closer examination, however, shows a number of problems:

- RRL only works for authoritative name servers; since most attacks do not directly use authoritative name

¹⁰<http://spoofer.cmand.org/summary.php>

¹¹A detailed discussion of RRL by its designer can be found at <http://ss.vix.su/~vixie/isc-tn-2012-1.txt>

servers but instead abuse open DNS resolvers (see also Sec. 2.1) its mitigating effect is limited. Especially if an attack uses queries that are cached by an open resolver, RRL on the authoritative name server for the domain abused for the attack will have little or no effect.

- As Paul Vixie's memorandum on RRL¹¹ mentions in Sec. 5, attackers can craft an attack that circumvents RRL by using a spread of queries in the attack rather than one single query.
- RRL can affect legitimate queries. Resolvers that do not (properly) cache results will suffer from RRL, which results in service degradation for clients behind these resolvers. Additionally, RRL can be turned into a denial-of-service weapon. By flooding authoritative name servers that use RRL with queries in which the source address is spoofed to be that of a legitimate resolver (e.g. the resolver of a large ISP) this legitimate resolver can experience service degradation in resolving specific queries on the authoritative name server under attack.

Thus, although RRL can certainly play a role in mitigation, it is not the definitive solution to name servers being abused for DNS amplification attacks.

6.3 EDNS0 cookies

Before EDNS0 was introduced the potential for amplification attacks using DNS was limited by the hard upper response size limit of 512 bytes. As we have shown in Sec. 2.1.2, EDNS0 is now in wide-spread use across the Internet. Because it allows for much larger responses, DNS amplification attacks have gained much more potential. As mentioned before, the root cause that allows for these attacks is source address spoofing. A potential solution to this is some form of source address authentication, that allows the recipient of a packet to establish that the source address has not been spoofed. An effective way to implement this is by using cookies, as proposed by Eastlake in [16]. In short, the idea is that a name server does not send large responses to a client using EDNS0 unless the client proves its authenticity using an authentication cookie established during an initial interaction between client and server. We believe that this can be a particularly effective solution to the DNS amplification attack problem.

6.4 Response size limiting

Another mitigation strategy is to limit the size of DNS responses such that no answer exceeds the acceptable upper limit for amplification we set in Sec. 3.2. The results of our analysis in Sec. 5 show that doing this would not affect most regular queries, but would probably only impact ANY queries. A beneficial side effect of this mitigation strategy is that it also prevents IP fragmentation of DNS responses, which is a serious stability problem affecting DNSSEC-signed domains, as described in [17]. If applied, DNS responses over UDP are restricted to a certain maximum size. Answers exceeding that size result in truncated DNS responses, forcing clients to retry the query over TCP. This effectively stops attackers from using responses larger than the chosen maximum size as TCP is connection-oriented and thus immune to source address spoofing.

6.5 Restricting or blocking ANY queries

The biggest amplification problem – if we do not consider crafted domains – is ANY queries. This begs the question whether this query type should be severely restricted or even blocked. Restricting ANY queries can be done using response size limiting (Sec. 6.4) and is automatically done by RRL (Sec. 6.2). A more extreme measure would be to block ANY queries altogether. This can only be done, however, if no legitimate use cases depend on ANY queries. We know, however, that certain software, such as qmail¹², uses ANY queries. Blocking ANY queries might thus affect some legitimate users.

Assessing the full impact of blocking ANY queries is outside the scope of this work but definitely warrants further investigation. We note that – just like RRL – blocking ANY queries is not a complete solution as other query types, such as DNSKEY queries, can still result in a significant amount of amplification.

7. RELATED WORK

Geva et al. [18] discuss DDoS attacks in a general sense. They performed simulations that illustrate the damage that attacks can cause and describe a spectrum of different DDoS attack mechanisms and various mitigation strategies against these attacks. They conclude that many defense mechanisms are problematic to deploy and that they may struggle to protect against the increasing threat level of today. A similar, but more practical and elaborate overview of DDoS attacks and defense mechanisms is given by Zargar et al. [19].

On the measurement side, Casalicchio et al. [20] describe a highly detailed reference architecture for measuring the stability and security of the DNS. They argue that the DNS is the most important infrastructure underpinning the Internet and that there is a great need to assess the health of the DNS on a continuous basis.

Rossov [21] performed measurements for 14 protocols that can be abused for amplification attacks. Where we focus on DNSSEC, Rossov covers amplification in its breadth. For DNS, Rossov considers amplification through open resolvers and directly abusing authoritative name servers. In the latter category he examines ANY queries for DNSSEC-signed domains. We corresponded with him and his data set covers 25K domains distributed over more than 70 TLDs. Rossov finds ANY amplification factors in the same range as we do, although the outliers in our data set exceed his findings and the distribution of amplification factors differs. We believe that this difference in distribution can be explained by the difference in sample size and composition; we use a much larger sample in which groups of domains hosted by a single operator may dominate. If such an operator has configured their servers such that they are less amenable to amplification then this will be reflected in the distribution of amplification factors in our results.

Highly relevant is work by Anagnostopoulos et al. [22]. They also perform extensive measurements to assess the potential impact of DNS amplification attacks. Where their work differs from our approach is that they focus on open resolvers. They measure numbers of open resolvers for three European countries and assess their capabilities (i.e. what amplification factor can be achieved by abusing them). This makes their work complementary to our study.

¹²<http://fanf.livejournal.com/122220.html>

8. CONCLUSIONS AND FUTURE WORK

We set out to answer the question how bad DNSSEC really makes DNS amplification. The simple conclusion from our measurements is that it is quite bad. If we only consider ANY queries, then DNSSEC-signed domains yield high amplification factors, averaging between 40 and 55. This exceeds the average amplification of regular DNS by a factor of 6×–12×. Looking deeper, however, the picture is more nuanced. For many common DNS queries, using DNSSEC results in larger responses but the amplification factor mostly stays within the acceptable upper limit (based on the maximum amplification of classic DNS) set in Sec. 3.2. Nevertheless, an attacker needs only one or a few domains with large amplification factors, and by carefully choosing a signed domain attackers can achieve significant amplification using e.g. DNSKEY queries. It is clear then that this needs to be addressed, both by mitigating the risk of carrying out a successful amplification attack as well as by improving the DNSSEC protocol, for instance by using cryptographic signature schemes with more favourable key and signature sizes that reduce the size of DNSSEC responses, such as elliptic curve digital signatures (ECDSA) [23]. We also bring back to mind that attacks with large amplification factors were already feasible without DNSSEC. Attackers already use crafted domains and can continue to do so regardless of whether DNSSEC is implemented. DNSSEC does, however, give attackers more options, such as directly abusing authoritative name servers instead of open resolvers and foregoing the use of crafted domains.

There are a number of countermeasures that can be deployed to mitigate the effect of DNS amplification attacks. The first two we discussed (BCP 38 and RRL) are already in active (albeit not universal) use. The other three, however, require further work. We consider EDNS0 cookies to be a particularly promising strategy for combating DNS amplification attacks but note that work on this approach is still at an early stage. We intend to investigate the effectiveness of this approach in a future project. Response size limiting, also requires more study. Before such a solution can become a recommended best practice the impact of this approach on legitimate DNS traffic will need to be assessed to ensure that it does not adversely affect DNS functionality. We have already looked at this on a small scale while working on [17] but a more extensive assessment is required. Blocking ANY queries, finally, also warrants further investigation as it may be a good stopgap measure against the worst amplification attacks. We note again, however, that it is certainly not a panacea as other query types can still lead to significant amplification.

As mentioned in Sec. 3.5, we will share the data sets resulting from our measurements as open data and do so on <http://traces.simpleweb.org>.

9. ACKNOWLEDGEMENTS

We would like to thank Si on Lloyd (Nominet, .uk), Marco Davids (SIDN, .nl) and Patrik Wallstr om (IIS, .se) for helping out with DNS data for their TLDs. Additionally, we thank VeriSign for sharing data on .com and .net and the Public Interest Registry (PIR) for sharing data on .org.

We would also like to thank Christian Rossov for reviewing the draft camera-ready version of our paper.

Finally, we would like to thank Shobha Venkataraman for shepherding our paper and the anonymous reviewers for their valuable feedback on our work.

Part of this work has been supported by the EU-FP7 FLAMINGO Network of Excellence Project (318488) and by the GigaPort3 programme funded by the Dutch Economic Structure Enhancing Fund (FES).

10. REFERENCES

- [1] L. Garber. Denial-of-service attacks rip the Internet. *Computer*, 33(4):12–17, April 2000.
- [2] D. Anstee, A. Cockburn, and G. Sockrider. Worldwide Infrastructure Security Report. Technical report, Burlington, MA, USA, 2014.
- [3] P. Mockapetris. RFC 1035 - Domain Names - Implementation and Specification, 1987.
- [4] D. Kaminsky. Black Ops 2008: It's the End of the Cache As We Know It. In *Black Hat USA*, 2008.
- [5] A. Cowperthwaite and A. Somayaji. The futility of DNSSEC. *5th Annual Symposium on Information Assurance (ASIA10)*, page 28, 2010.
- [6] D.J. Bernstein. High-speed high-security cryptography: encrypting and authenticating the whole Internet. In *27th Chaos Communication Congress (27C3)*, Berlin, 2010.
- [7] M. Kühner, T. Hupperich, C. Rossow, and T. Holz. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *Proc. of the 23rd USENIX Security Symposium*, August 2014.
- [8] J. Damas, M. Graff, and P. Vixie. RFC 6891 - Extension Mechanisms for DNS (EDNS(0)), 2013.
- [9] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzer: Illuminating the Edge Network. In *Proceedings of IMC 2010*, pages 246–259, New York, USA, 2010. ACM Press.
- [10] R. Vaughn and G. Evron. DNS Amplification Attacks (preliminary release). 2006.
- [11] J.J. Santanna and A. Sperotto. Characterizing and Mitigating The DDoS-as-a-Service Phenomenon. In *8th International Conference on Autonomous Infrastructure, Management and Security (AIMS)*, LNCS, Brno, Czech Republic, 2014. Springer.
- [12] O. Kolkman and R. Gieben. RFC 4641 - DNSSEC Operational Practices, 2006.
- [13] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. RFC 4034 - Resource Records for the DNS Security Extensions, 2005.
- [14] O. Kolkman, W. Mekking, and R. Gieben. RFC 6781 - DNSSEC Operational Practices, Version 2, 2012.
- [15] P. Ferguson and D. Senie. BCP 38 - Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing, 2000.
- [16] D. Eastlake. Domain Name System (DNS) Cookies (Internet Draft), 2014.
- [17] G. van den Broek, R. van Rijswijk-Deij, A. Sperotto, and A. Pras. DNSSEC Meets Real World: Dealing with Unreachability Caused by Fragmentation. *IEEE Communications Magazine*, 52(4):154–160, 2014.
- [18] M. Geva, A. Herzberg, and Y. Gev. Bandwidth Distributed Denial of Service: Attacks and Defenses. *IEEE Security & Privacy*, 2013.
- [19] S.T. Zargar, J. Joshi, and D. Tipper. A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks. *IEEE Communications Surveys & Tutorials*, 15(4):2046–2069, January 2013.
- [20] E. Casalicchio, M. Caselli, and A. Coletta. Measuring the global domain name system. *IEEE Network*, 27(1):25–31, January 2013.
- [21] C. Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Proceedings of the 2014 Network and Distributed Systems Security Symposium (NDSS 2014)*, number February, pages 23–26, San Diego, 2014. Internet Society.
- [22] M. Anagnostopoulos, G. Kambourakis, P. Kopanos, G. Louloudakis, and S. Gritzalis. DNS amplification attack revisited. *Computers & Security*, 39(February 2012):475–485, 2013.
- [23] P. Hoffman and W.C.A. Wijngaards. RFC 6605 - Elliptic Curve Digital Signature Algorithm (DSA) for DNSSEC, 2012.