.se Health Status

– DNS and DNSSEC

1 About .SE's Health status focus area

This survey is included in of .SE's Health status focus area. The aim of this focus area is, among other things:

- To monitor the quality of the Internet's infrastructure in Sweden by compiling and analyzing facts,
- To disseminate the results from the surveys, and
- To use advice and recommendations to contribute to ensuring that the infrastructure functions well and has a high level of accessibility.

Another aim is to, when necessary, detect deficiencies and improprieties.

The Health status report was financed by .SE. Results from the study have been processed and analyzed, and the report was compiled by Anne-Marie Eklund Löwinder, Quality and Security Manager at .SE. In-depth analysis of DNSSEC was conducted by Patrik Wallström, Project Manager at .SE. Patrik also has operational responsibility for the tools that were used.

A special thanks to Fredrik Ljunggren, Kirei, Patrik Fältström, Netnod and Torbjörn Eklöv, Interlan for valuable insights.

More information about the content of the report is available from Anne-Marie Eklund Löwinder. Her e-mail address is anne-marie.eklund-lowinder@iis.se. More information about the technology behind this study can be obtained from Patrik Wallström. He can be reached at patrik.wallstrom@iis.se.

2 Introduction

The purpose behind Health status is to make registrants, especially those of important functions in society and their service providers, aware of deficiencies and improprieties by presenting reports with results and analysis from tests that we perform with some degree of regularity.

In order to effectively be able to influence the quality of the Swedish Internet infrastructure, we develop and maintain our own testing tools and use others that we package into a common platform. We are continually working to improve the tools, the analysis of results and reporting.

For the past five years, we have published an annual compilation report that shows the results of our surveys, where we largely strive to conduct a follow-up on the previous year's survey in order to detect patterns of development and trends.

For 2012, we have partially complemented the strategic focus of the Health status with the aim of being able to present more in-depth analyses in report form every quarter, based on more in-depth trials within certain specific areas.

The most recent report is the result of an in-depth study and analysis within the DNS and DNSSEC areas that was conducted in February 2012.

The goal of the survey is to chart and analyze the quality and reachability of the Domain Name System (DNS) in the .se zone, for a selection of domains that represent important functions in society, all signed .se domains and a random selection corresponding to 1 percent of all .se domains.

The report is primarily geared toward those responsible for the operation and management of an entity's IT and information systems.

2.1 Purpose of the survey

The purpose of focusing on and studying DNS and DNSSEC more in-depth is to draw attention to the problems and deficiencies that quite a few domains in the .se zone are suffering from.

2.2 About the survey group

In the current survey, a **survey group** of a total of 911 domains allocated across 1,374 unique name servers, both IPv4 (1,241 domains) and IPv6 (133 domains), were tested. The term "unique" is defined as servers with unique IP addresses. A name server with a service provider can, of course, house several domains.

In addition, a comparison was made with the entire .se zone, based on a **control group** that consists of 1 percent of the entire .se zone, or 11,730 randomly selected .se domains, allocated across 2,938 unique name servers – 2,698 using IPv4 and 240 using IPv6. Finally, we studied **all signed domains in the .se zone** to see how well DNSSEC is set up and operating.

During the survey period, approximately 14 percent of all active domains in the .se zone were signed using DNSSEC.

2.3 Limitations

It was our intent to also conduct a comparison of the quality of the DNS between unsigned and signed domains. Such a comparison proved to be irrelevant, however, since a very small number of name server operators account for a majority of the signed domain names in .se. Those who have signed domains with DNSSEC are some of the largest name server operators, which already have a relatively good infrastructure, and it would be deceptive to compare the quality of those with the other operators.

It is worth noting that .SE's three largest registrars account for 50 percent of the market, while the seven largest commands 75 percent. Among the name server operators, the two largest have 36 percent of the market, while the five largest commands 50 percent. Also, among the name server operators, there are a vast number of very small players.

One limitation in terms of the selection of the control group, is that the previous version of the survey platform was able to generate a list of an exact number of randomly chosen domains and retain that over time. Now the tool is following a different type of algorithm making it impossible to determine the exact number of subjects for the survey group or which ones those should be. Instead we select a sample, in this case 1 percent, which for the actual trial amounted to 11,730 domains. As a result, the control group will change with each survey.

2.4 Test subjects

Tests were performed on domains and name servers for a large number of important public organizations; 911 domains in total, grouped into the following categories:

- 60 public service and state-owned companies
- 79 banks, financial institutions and insurance companies
- 22 Internet service providers (ISPs)
- 290 municipalities
- 21 county councils
- 34 media companies
- 228 government agencies and public service companies, including county administrative boards
- 39 universities and colleges
- 146 registrars

In addition to this list, we also had a control group with 11,730 randomly selected domains.

3 Summary

The survey was conducted during February 2012 and focused on the quality of DNS and DNSSEC. The development of DNSSEC and IPv6 are important parameters, which is why it was also highlighted in the government's strategy for the IT-political area, "IT in the public service – a digital agenda for Sweden."[1]

Our recurring survey of DNS for operations that we deem critical showed small changes compared with the survey conducted in autumn 2011. Our control group, which is representative of the .se zone as a whole, did show more serious defects now than during the previous survey – 22.5 percent of the name servers had serious defects now compared with 17 percent in October (see also Section 5.1).

The percentage of open recursive name servers remained unchanged compared with the autumn 2011 survey, which concerns us. There are few reasons to operate an open recursive name server, and at the same time, they comprise a security threat since they can be utilized by an attacker for targeted Denial of Service attacks. Consequently, we would prefer that no open recursive name servers at all existed in the .se zone. (For more information about the problem with recursive name servers, see Sections 5.2 and 5.5.)

In the new portion of our survey, meaning the testing of all DNSSEC signed .se domains, the overall results were approved. When the study was conducted in February, there were over 170,000 signed domains in the .se zone, of which over 150,000 were signed during .SE's DNSSEC campaign in December 2011. It is apparent that the name server operators generally addressed the early problems they experienced.

Almost all of the domains use Key Signing Keys (KSK), which are 2,048 bits long and Zone Signing Keys (ZSK), which are 1,024 bits. Those are satisfactory values, although we anticipate being able to use somewhat longer keys in the future. In our opinion, too many domains are using 512 bit long keys, which we consider to be a weak choice for key length in 2012.

One administrative problem is that many DNSSEC signed domains appear to be lacking a link between the period of validity for the zone and the period of validity for their signing keys – in other words, the value in the field *SOA Expire* often lacks a link to *RRSIG expiration* time. This may result in problems with reachability since you run the risk of keys becoming invalid without warning. Another problem is domains that have signature lifetimes that are unusually long or too short. We would definitely encourage general improvements in that area.

Our concluding recommendation is that name server operators in the .se zone should transition to only publishing DS records of the new more secure type 2.

See Section 6 for a more in-depth discussion regarding the DNSSEC portion of the survey.

[1] http://www.regeringen.se/content/1/c6/18/18/01/509f1b0c.pdf

4 The Domain Name System, DNS

A very important function of the Internet's infrastructure is the Domain Name System (DNS). DNS is a globally distributed database with domain names connected to one or more IP addresses stored on millions of servers. Every piece of Internet-connected equipment is identified using a unique IP address, which is used to convey data packets on the Internet to the correct location.

DNS is the system that allows us to "surf" the Internet in a user-friendly fashion. DNS is also, practically speaking, a prerequisite for being able to send and receive e-mail. A major advantage of using domain names is that the corresponding IP address can be changed without affecting the domain name.

Activities that are dependent on being accessible via the Internet need to protect accessibility to the name servers, since those are what point out where their own resources reside. Maintenance and administration of name servers can either be done by an external DNS operator or in the capacity of one's own internal operations.

Those who oversee the maintenance of their own name servers can increase accessibility by outsourcing the operation of a secondary name server to one or more external DNS operators. For an even more robust solution, it is important that name servers be situated within the networks of various Internet service providers in order to ensure that DNS works even if one of the operators has a problem.

Resolvers are ind of type of name server that have a special function. These are generally operated by Internet service providers and provide a look-up service of domain names when requested by individual computers, for example, translating www.iis.se to an IP address and vice versa.

All name servers need to have an updated, modern version of the software with the latest updates in order to be protected against logical vulnerabilities.

A registrant can check if their domain is correctly configured by using .SE's tool DNSCheck[2].

4.1 High quality DNS

The survey results we have seen over the years reinforce our hypothesis that there is a general lack of knowledge regarding what is required to maintain a high quality Domain Name System (DNS), even if the definition of "high quality" is always debatable. There is also reason to believe that this lack of knowledge is demonstrated through deficiencies both in terms of operation and operational responsibility.

[2] http://dnscheck.iis.se/

In this case, we personally defined what we believe to be high quality, although we used the recommended international industry standard, *Best Common Practice*, as the basis for our definition.

Questions that we attempt to answer:

- How does the organization manage its DNS?
- Who is responsible for DNS within the organization, what is its structure (in relation to what can be considered to be industry standard or Best Common Practice, BCP)?
- What are the most serious deficiencies?
- In what categories do they most frequently occur?

The data-collection process was automated and included testing of the most frequently occurring errors and deficiencies we associate with DNS operation compared with what is considered common practice. Based on these tests, we investigated how well the organizations' systems function in various contexts, the areas in which the most serious defects arise and the possible consequences. We have also linked this information to general recommendations on how we believe the Swedish DNS infrastructure should be structured.

With this survey we have also conducted a number of DNSSEC-specific measurements following the extensive increase in DNSSEC-signed domains, following a campaign that was conducted in December 2011.

4.2 .SE's contribution to improving the Swedish DNS infrastructure

Our surveys and reports are one way to inform the general public and those responsible at various entities of the need for improvement measures. We also naturally work in close contact with our registrars.

We also operate the publically accessible service DNSCheck, which is mentioned in Section 4 above. The purpose is to offer tools that provide assistance for self-help through indications about problems that may need to be addressed.

By working with strategic partners such as the Swedish Post and Telecom Agency (PTS), the Swedish Civil Contingencies Agency (MSB), and the Swedish Association of Local Authorities and Regions (SALAR) .SE is working to ensure that municipalities, through county councils, are able to apply for grants to conduct projects to implement DNSSEC. We would also like to see agencies and individuals in positions of authority use our advice and recommendations and apply suitable measures for improvements with the areas of DNS, DNSSEC and IPv6.

5 DNS-related testing

Outlined below are the results of the in-depth DNS test. Just as before, we are reporting two different types of DNS-related problems and categorizing them as errors and warnings respectively.

Error: Anything marked as an errorin the study should be considered serious and something that could immediately impact operations negatively and should be corrected as soon as possible so that the organization can be assured of a high level of availability and accessibility in DNS and other resources. The most common errors are accounted for below.

Warnings: Warnings also constitute errors that could affect operations, although they are not as serious and corrective actions are not deemed as urgent, quality and availability would naturally be enhanced if those defects were eliminated.

5.1 Serious errors

Survey group (911 domains)
-1.5 percent

Control group (11,730 domains)
+5.5 percent

Compared with results for the survey group during the annual test in 2011, the results from this testing were more positive, with 19.5 percent serious errors compared with 21 percent from the October test.

For the control group, the entire .se zone, the results were worse – 22.5 percent serious errors compared with 17 percent in autumn 2011.

In other words, the total number of serious errors fell slightly in the survey group, while they increased in the control group.

Of the domains in the survey group, 52 domains had one serious error, or approximately 6 percent. Some 72 domains had two serious errors and 54 domains had three or more errors.

Of the 11,730 domains in the control group, the corresponding figures were 1,237 domains with one serious error, e.g. slightly more than 10 percent, 815 with two errors and 591 domains with three or more errors.

Part of the changes in the results compared to the previous test can be explained by the fact that the control group changes from one time to the next since a new sampling, comprising 1 percent of randomly selected domains, are selected from the zone file for every survey. At the same time, however, the control group corresponds to a representative sample.

The two domains that had such serious errors during last year's major survey that they were rendered impossible to test, still "live" on. According to our tests with DNSCheck, they are not even in the domain name system, yet it is still possible to surf to their websites. It is probably impossible to reach them via e-mail. The domains are likely used just for web traffic and not for e-mail, which is probably the reason that the registrants do not notice that they have problems with accessibility to their domains.

The problem with both of these domains is that there are so-called "glue-records" in the .se zone that points to two name servers. Of those two servers, one gives a REFUSED message to all domain inquires, while the other response with CNAME to all inquiries. That CNAME is linked to a name that only has one A record. The effect is that the web browser that attempts to find an IP address to connect with actually gets one, but nothing else works. Our tool, DNSCheck, requires there to at least be a designated name server that responds appropriately to NS queries to the domain, which is lacking for both of these.

Nevertheless, it is a strange phenomenon that demonstrates what we usually say about DNS – that it is extremely forgiving and will continue to work even if someone does something very wrong.

5.2 Most frequently occurring errors

Among the domains and name servers tested, the most common errors were:

- The name server did not respond to queries via the TCP (Transmission Control Protocol). This is probably because DNS server was not correctly set up or a firewall was incorrectly configured. It is a fairly common misconception that DNS does not need to communicate according to the TCP protocol (if it does not provide zone transmissions).
  However, TCP is usually a requirement under a standard (RFC 5966, *DNS transport over TCP implementation requirements*), and the trend is that the need for TCP is increasing as new protocols result in it being used more extensively than in the past. This defect indicates that the person who configured the name server has insufficient current knowledge of DNS.

- The organization has an inconsistent name server infrastructure (NS). The name servers listed with NS records in a child zone differ from the information found in DNS in the parent zone and, accordingly, the name servers cannot assume authoritative and proper responsibility for the domain. If the information is not consistent, the availability of the domain is negatively affected, which indicates deficiencies in the internal DNS management. Some examples of such inconsistencies are provided below:
  - The IP address of a name server in the child zone is not the same as in the parent zone in the level above. This is a configuration error and should be corrected as soon as possible. The administrator of the domain has probably forgotten to perform an update after a change was made.
  - A name server is listed in the parent zone but not in the child zone. This is probably an administrative error. The parent zone must be updated as soon as possible so that it lists the same name servers as those listed in the child zone. The consequence of such an error is that the redundancy that someone has tried to create essentially does not exist.

- The name server lacked EDNS support. This is an expansion of the DNS protocol to handle DNS responses that exceed the UDP protocol's limit of 512 bytes. EDNS enables DNS responses in excess of this amount, which is also becoming increasingly normal along with the expanded use of DNS in conjunction with, for example, DNSSEC and IPv6.

- DNS server did not respond to orders via UDP (User Datagram Protocol). The probable reason is that DNS server was not correctly set up or the firewall was incorrectly configured. Since a name server that responds to neither TCP nor UDP is probably not reachable at all, the error may be found elsewhere, for example in connection with the name server, or the server may not have a correctly stated IP address. Name server tests are now finalized if both of these conditions have been confirmed.

- Only one DNS server is found for the domain. There should always be at least two DNS servers for a domain so that temporary errors with connections can be handled. If one of the servers or the connection to it were to stop functioning, services directed from the name server would also be rendered unavailable. We made separate calculations for IPv4 and IPv6. We consider that having an insufficient number of servers is a more serious problem for IPv4 (causes errors), while we currently consider it a less serious problem for IPv6 (generates a notification).

- The name server is recursive. The name server responds to recursive requests from third parties (as in DNSCheck). It is very easy to abuse open recursive resolvers during distributed denial of service attacks (DDOS), since the use of a very small DNS query can create an amplification effect generating exponentially larger responses. False sender addresses can be generated using DNS, and those who want to attack a system create queries under a false sender address that produce major DNS responses that are sent to the presumed sender, which is in fact a third party whose services can more or less be blocked (refer to Section 5.5).

- The start of authority (SOA) serial number is not the same in all name servers. This is usually due to an incorrect configuration, but is sometimes due to slow distribution of the zone to secondary DNS servers. This means that users searching for resources under a domain may receive different responses depending on which name server receives the request, since the name server would then contain differing information on domains.

- The name server yields a SERVFAIL when contacted. SERVFAIL is the response one gets when DNSSEC does not work for a domain, but also if something on the server side is incorrectly configured or if the name server has other problems handling a query.

5.3 Warnings

Survey group (911 domains)
+/-0

Control group (11,730 domains)
+13 percent

The percentage of warnings in the survey group was 37 percent, which was unchanged since last autumn's test. In the control group, the number of warnings increased from 45 to 58 percent. In other words, only 4,912 domains passed the control without warnings, while 187 domains generated one warning, 186 domains generated two warnings and as many as 6,443 domains generated three or more warnings. The most common warnings are e-mail related because people who set up zone files do not record real e-mail addresses in the DNS SOA rname (responsible name) field due to the risk of receiving large amounts of spam.

Many warnings are also generated as a subsequent result of something else being reported as an error. In many instances, several warnings are generated from the same actual error. If you fix the error then you fix a number of warnings.

5.4 Providers of name service operations

Normally, it is the registrar who both administers a domain and is responsible for operating a name server for that domain. .SE's seven largest registrars together handle 75 percent of the domains in the .se zone. It is highly probable that serious faulty configurations among registrars who also operate name servers for their customers would be very noticeable.

5.5 Recursive name servers

Survey group (911 domains)
+/-0 percent

Control group (11,730 domains)
+/- 0 percent

The percentage of name servers open for recursion dropped during the 2011 survey and was down to 11 percent compared with 15 percent in 2010. During the current survey, that percentage remained unchanged since the last time we tested. The most commonplace occurrence is still among municipalities. That is

a phenomenon that we would prefer to see disappear entirely, for the reasons listed below.

A **recursive name server** not only responds to queries about DNS records for which it itself is responsible, but also goes further and asks other name servers to respond to queries. These queries can be both labor-intensive (meaning that they utilize extensive computer capacity) and result in a relatively large amount of data, which means that organizations normally want to limit the number of persons permitted to use the recursion function.

An **open recursive name server** responds to all queries it receives for which recursion has been requested. This makes it possible for external parties to launch Denial of Service attacks; for example, via the open name server by allowing these parties to submit queries that will result in unusually large responses (what is known as an Amplification Attack). Combined with a false sender address that leads to the response being sent somewhere else, this comprises a Denial of Service attack.

As we have pointed out every year since 2007, open recursive name servers have very few legitimate areas of application and can be abused in conjunction with denial of service attacks, among others. It is therefore strongly recommended to eliminate the possibility of utilizing open recursive resolvers with the assistance of available techniques for that purpose.

Today, name servers are delivered with recursion turned off as the default setting. We would also like to think that those responsible for the DNS infrastructure have become better at implementing a separation between authoritative name servers (those that actually should respond to queries) and resolvers (those that just convey queries and responses).

The fundamental problem is not actually open recursive name servers, but the fact that Internet service providers do not filter traffic by sender addresses. If they did, open recursive resolvers might not be considered a problem.

Unfortunately the service providers are reluctant to implement such filtering. Accordingly we need to attempt to limit the damage caused by DDOS attacks until the service providers have identified cost-effective solutions to the fundamental problem.

Closing a recursive resolver is a relatively simple task that is worth the trouble of implementing, since it will help ease problems arising from DDOS attacks.

Below, we have gathered some links to high-quality, informative material about DDOS and open recursive name servers and how they can be counteracted.

Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing
http://tools.ietf.org/html/bcp38

Secure Domain Name System (DNS) Deployment Guide
http://csrc.nist.gov/publications/nistpubs/800-81r1/sp-800-81r1.pdf

DNS Amplification attacks
An excellent description of how these attacks occur and what they entail.
http://www.isotf.org/news/DNS-Amplification-Attacks.pdf

Official advice from the US CERT

The Continuing Denial of Service Threat Posed by DNS Recursion
http://www.us-cert.gov/reading_room/DNS-recursion033006.pdf

ISC BIND. Here you can find source codes and binaries for BIND and links to highly interesting and useful information.
https://www.isc.org/downloads/all/

BIND 9 Administrator Reference Manual.
Includes examples of configuration, practical tips and detailed descriptions of BIND functions.
http://oldwww.isc.org/index.pl?/sw/bind/arm95/index.php

6 DNSSEC

DNSSEC protects Internet users from forged or manipulated DNS information, for example, what is known as DNS cache poisoning. Responses to DNS queries that are protected using DNSSEC are assigned a digital signature, the verification of which ensures that DNS information has not changed en route from the name server to the recipient system.

In the following sections we report how many domains in the .se zone are signed using DNSSEC and how they are set up. We also explain which parameters are important to keep an eye on and describe why it is important.

6.1 Usage of DNSSEC in the survey group

Among the domains in the 2011 survey group, 6.69 percent, or 61 domains, were signed using DNSSEC. Municipalities, government agencies, county councils and ISPs are the primary organizations that have begun to implement the safer technology.

Due to an extensive campaign, we saw a very large increase in DNSSEC signed domains in December 2011.

For the current survey conducted, the number of signed domains in the survey group increased to 81 domains, corresponding to 8.89 percent. The control group, during this same survey, recorded 1,556 signed domains, or 13.37 percent. This can be compared with the autumn study when only 0.45 percent or a total of just 50 domains in the control group was signed.

6.2 DNSSEC specific testing

After some startup issues (DNS operators that had problems with their signed domains), we are starting to take a closer look at how DNS operators actually signed their domains.

Above all, we feel it is important to give early warning about domains that do not work. We would also like to get an indication and a warning about which domains are within the risk zone of losing their accessibility, which is why we have also looked at such parameters in child zones (secondary domains under .se) that are associated with DNSSEC. This involves, among other things, the lifeline of DNSSEC signatures, and how big of a time margin there is before they expire and the domain ceases to function.

DNSCheck is not quite fully developed to handle these values, which is why we developed a new tool that only looks at DNSSEC for our signed domains.[3] Below we primarily present results from these measurements. These readings were also taken in February 2012.

The judgments we make for the selection of suitable values for DNSSEC are based on the content in RFC4641bis, which, at the time of writing, is the subject of discussion within the IETF and not yet published as formal RFC.[4]

[3] https://github.com/dotse/dnssec-analysis
[4] http://tools.ietf.org/html/draft-ietf-dnsop-rfc4641bis-09

6.3 Works, does not work

The total number of DNSSEC signed domains for the entire .se zone was, at the time of analysis (February 10), 174,487 out of a total of 1,195,719 registered domains (approximately 50,000 of those lacked delegation, meaning that they were registered but not in use, which is why there is no DNS quality to measure).

The error messages you receive when DNSSEC does not work for a domain is SERVFAIL, which is unfortunately the same message received when something on the server side is not correctly configured, or if the name server software has other problems handling the query. That makes it less than easy to determine whether an error is due to DNSSEC.

By DNSSEC signed domains, we mean domains that have a DS record publicized in the .se zone, which means that the zone must work with DNSSEC turned on. A DS record must match a published DNSKEY in the zone, which in turn signs over the records that are published in the zone.

Of the total of 174,487 signed domains, there were 163,700 functioning ones, the remaining 10,787 domains (6 percent) did not work at all when the name servers with DNSSEC turned on on the resolver side attempted to verify signatures. It is obvious that there was a major learning curve during .SE's DNSSEC campaign last December, although those problems have been sorted out following the conclusion of the campaign. We have chosen to only look at DNSSEC parameters for domains that have functional DNSSEC. With the large number of resolver-operators that are using DNSSEC today, we expect that it will become noticeable very quickly for a registrant if something is wrong.

Since our testing method poses five different DNS queries (A, DNSKEY, MX, NSEC3PARAM and SOA) to the authoritative name server that is handling the domain, we receive a few more SERVFAIL calculated according to the number of domains in a few more cases than 10,787, namely on 10,808 different name servers. In short, that means that, for example, an older implementation of a name server that cannot respond to a query for NSEC3PARAM results in SERVFAIL. We can see this a little more clearly if we examine SERVFAIL by query type:

RR type	Number
A	10 793
DNSKEY	10 787
MX	10 789
NSEC3PARAM	10 792
SOA	10 791

The exact reason as to why there is a small difference when we look at various record types (RR type) is somewhat unclear, but it is likely due to either outdated software or customized implementations of DNS protocol that do not handle all record types correctly, meaning by standard method.

If we examine the domains that do not function in greater detail, about 20 percent are under the largest DNS operator (of those who are using DNSSEC). After that, there is a very long list of name servers that have domains that have a DNSSEC signed domain linked to it. Altogether there are 8,390 name servers on the list. If you look at the total number of domains (even non-signed) that we have a total of 47,606 name servers that have at least one domain delegated to it, which yields over 17 percent of name servers with support for DNSSEC. Note, however, that the total number of name servers is not the same as the number of DNS operators. A DNS operator usually uses at least two name servers.

What is most critical for a functional DNSSEC signed zone is that there is a DS that matches the DNSKEY in the child zone, and that it has valid signatures. That is why we have also taken a closer look at signatures in particular.

6.4 Signature lifetimes

Signatures are normally used over all records in a zone (all authoritative data is signed), but also over certain delegation data (depending on if it is NSEC, NSEC3 opt-in or NSEC3 opt-out). Here, we are looking at signatures in three different ways: *inception time* is when the signature is created, *expiration time* is when the signature expires. We also take a look at how SOA Expire is related to *expiration time*, since these two parameters are intimately connected. The difference between *expiration time* and *inception time* is the signature's total period of validity.

Since we have measured DNSSEC through a snapshot view of the zone, and not over a longer period of time, we do not see the frequency with which signatures are updated. Consequently, signature lifetimes can vary depending on when the test was conducted. For future tests, we could take daily measurements over a period of a couple of weeks to also see these types of updates.

In order to simplify the presentation of results, we have divided the signature lifetimes into number of days. The graph below of *inception time* shows the age of the signatures that we saw. Shown below are the average values for the five questions we see signatures from. We were unable to determine whether the average, smallest value or largest value differ noticeably, and are therefore displaying average values. All of the graphs of the signatures below display the number of **domains on the Y axis**, and **the number of days on the X axis**.

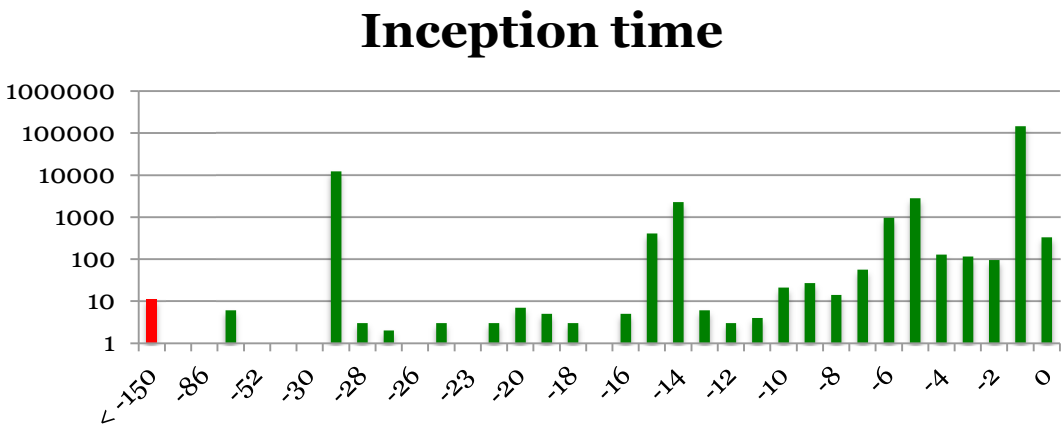

The graph above shows that most of the signatures are between a couple of weeks to a couple of days old. Nor are one-month old signatures a big problem, but if we look at the red bar we can see the extremes. Having signature lifetimes of a half-year is not necessarily favorable. That opens up for the possibility of replay attacks[5] for those who would send out erroneous data for something that perhaps should not exist or that has changed.

A signature that is created the same day or the day before the test is still indicative of regular, perhaps daily, updating of signatures, which is good.

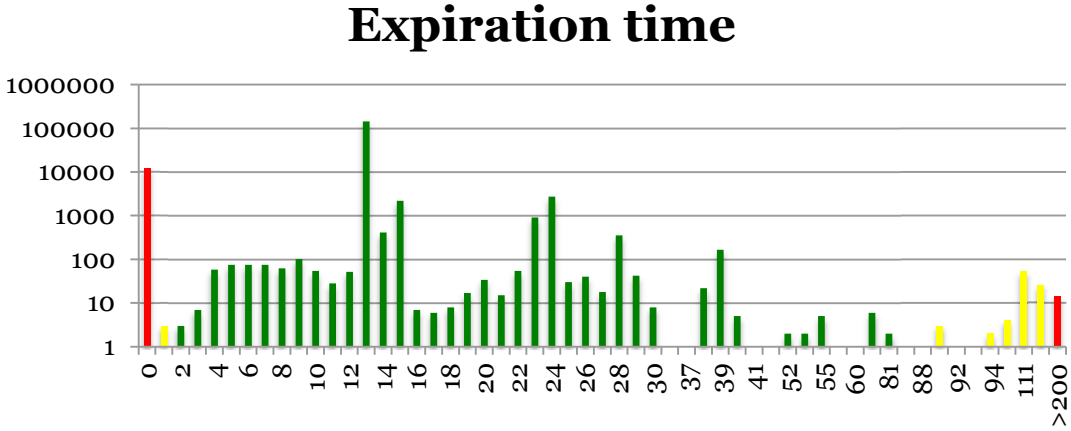

If we instead look at *expiration time* – meaning the time from which a DNSSEC signature becomes valid up until it becomes invalid – we can see that a considerable number of signatures, slightly more than 10,000, expire the same day that we conducted the test (the large red bar). That means that there is the risk of major problems if the equipment that signs the domain breaks down, or if you simply lose your Internet connection.

If, on the other hand, you have validity periods that are too long, as we see on the right side of the graph (the bars highlighted in yellow or red), then you risk being subjected to replay attacks. A reasonable balance between the operational

[5] http://en.wikipedia.org/wiki/Replay_attack

necessity to be able to handle crises (such as broken equipment, being disconnected from the Internet, or stolen keys) and the risk of replay attacks needs to be made, and that assessment should be made according to what the redundancy looks like in relation to the risk of attacks against DNSSEC signed data. You should also keep in mind how often signatures are renewed, which should be done frequently, regularly and preferably using automation.

Having long periods of validity for signatures can be just fine, if you conduct key rollover at shorter intervals. However, that also means that you just as routinely replace the corresponding DS records in the parent zone, otherwise you still runs the risk of replay attacks. (Simply replacing ZSK will not help, since such a key can also be subjected to a replay attack.)

It is not inconceivable that products (through automation) will implement this strategy in the future.

6.4.1 SOA Expire vs. RRSIG expiration

The expire value of SOA records are there to ensure the stability of the zone. A zone is served by a number of secondary name servers, and in those instances when the main server is unreachable for a longer period of time, it is appropriate that the secondary name server also stops responding to queries about the domain. That is also one way to make DNS operators aware of this sort of network problem.

RIPE's recommendation for an appropriate value for SOA Expire is 1,000 hours, which corresponds to approximately 41 days.[6] Since we have a new parameter in the DNSSEC that controls the zone's validity, namely RRSIG expiration, we should think about the role of the SOA Expire parameter. In RFC4641bis, the idea is that SOA Expire should be about two-thirds longer than RRSIG expiration. The reason for this is that one does not want to risk having signatures that expire. It is more appropriate to make the DNS operator aware that the secondary name server has not received an updated zone at an earlier stage than usual. Accordingly, we also examined the relationship between SOA Expire and RRSIG expiration in further detail.

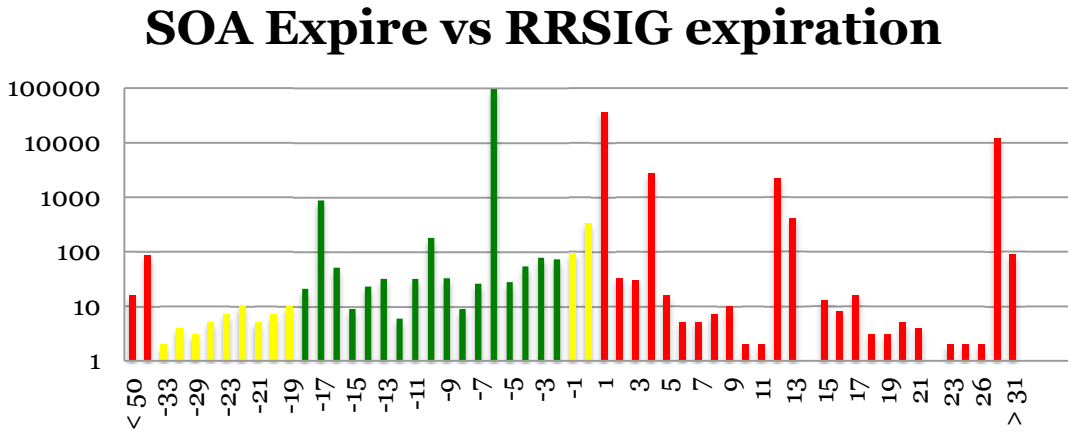

[6] http://www.ripe.net/ripe/docs/ripe-203

Without making any judgment of the actual condition at two-thirds, we examined the difference between SOA Expire and RRSIG expiration for a number of days, which is what the graph above shows. To the right of zero, are the zones that have a higher value for SOA Expire than for RRSIG expiration, which means that the signatures expire before the name server stops handling the domain. However, in both cases, this means that a resolver will receive SERVFAIL from the secondary name server instead of a correct reply. The green bars in the graph represent reasonable value, yellow are slightly more questionable, and we consider the values presented in red as unsuitable.

Those who chose to sign using DNSSEC will need to make a determination between how fast one would like to find out about errors, and how fast one can rectify them. In the event that secondary name servers cannot retrieve a new zone profile from the primary name server, it means that the secondary name server runs the risk of not receiving the revised data, despite having the correct signature for the existing zone content. Consequently, it is important to identify problems with the secondary name servers quickly.

This can be achieved in part through monitoring of name servers by looking to see that one is getting the correct DNS reply, and in part by looking up the serial number for the zone (which is also available in the SOA record) to check that all name servers have the same version of the content in the zone.

On the left-hand side of the graph we also see a significant number of domains that have a considerable margin in the other direction, meaning an SOA Expire that is significantly longer than a RRSIG expiration. We suspect that this is not a conscious decision by the DNS operator. It may be worth examining these parameters an extra time and to reassess your DNS risk management.

For the signed domains, SOA Expire is as follows:

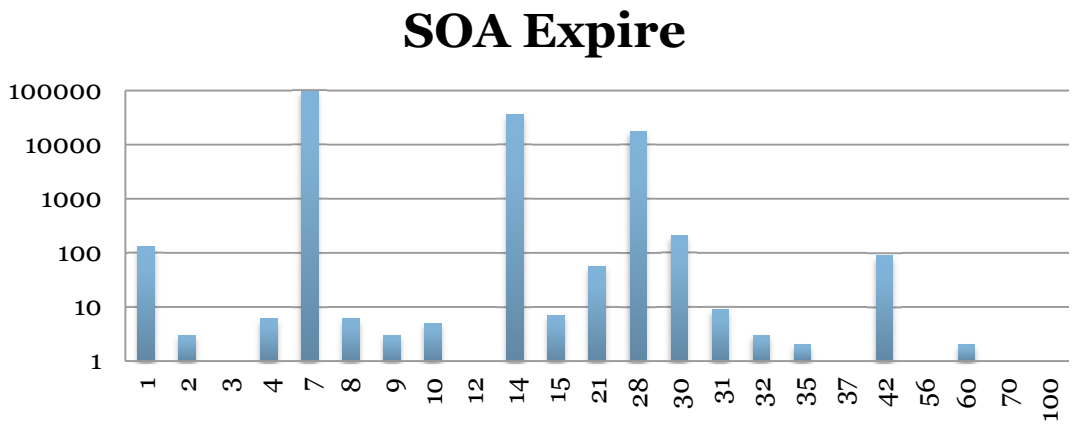

The most commonplace is an SOA Expire with a value from one week up to one month. During testing, however, we did not receive any SOA record from 11,212 zones. The fact that we did not receive SOA from these zones could be due to those name servers using different methods for zone distribution than the usual (meaning AXFR, IXFR or dynamic updating) and that use information in SOA records. An SOA Expire for one day is not very appropriate.

6.5 DNSSEC keys

We have also chosen to look more closely at the keys that are used for DNSSEC. There are a number of choices to make when it comes to keys and the administration of DNSSEC. In addition to the actual storage (it is important to store the private keys as securely as possible) it is also possible, depending on security level requirements, to chose different key lengths and algorithms. The oldest defined key types that are predefined in the standard for DNSSEC are not compatible with NSEC3, for example.[7]

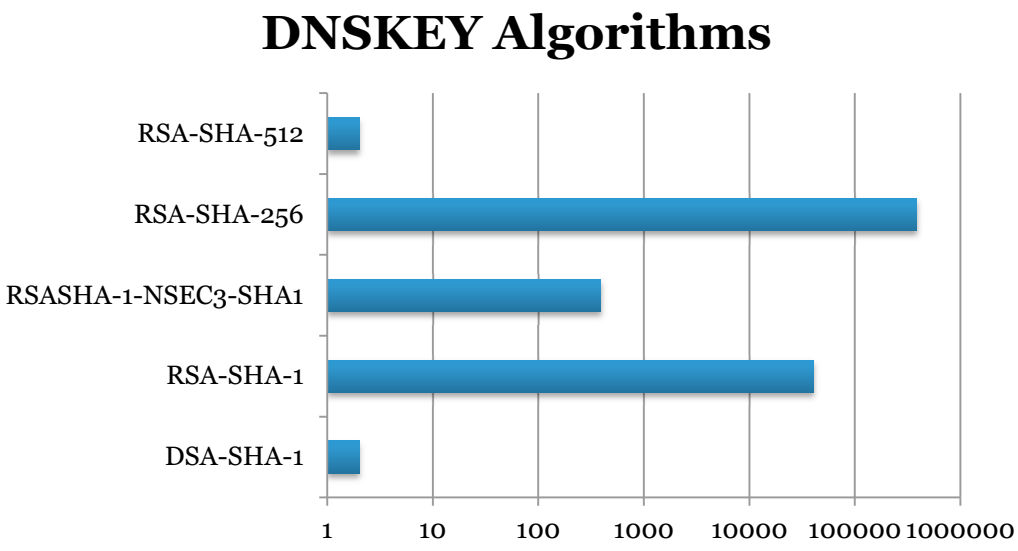

The most common algorithm is one of the newer ones, namely RSA-SHA256. These keys generate RSA signatures with SHA256-hashes, which are considered more secure than SHA-1. Within the framework for DNSSEC, this means that it is significantly more difficult to create hashes that collide with other information. Security for hashes is dependent on having two different amounts of data not happen to generate the same hash.

In the graph above we also find two keys with the algorithm DSA-SHA1. There is just one domain that uses this algorithm, allocated to KSK (Key Signing Key) and ZSK (Zone Signing Key). We advise against the use of DSA for DNSSEC since it requires far more CPU to validate signatures made with DSA, and also requires a good selection of random numbers when generating signatures.

We can also see how keys are used by looking at the algorithms that exist in the signatures we find, and that corresponds to the keys that are available:

[7] RFC 4033, 4034 och 4035.

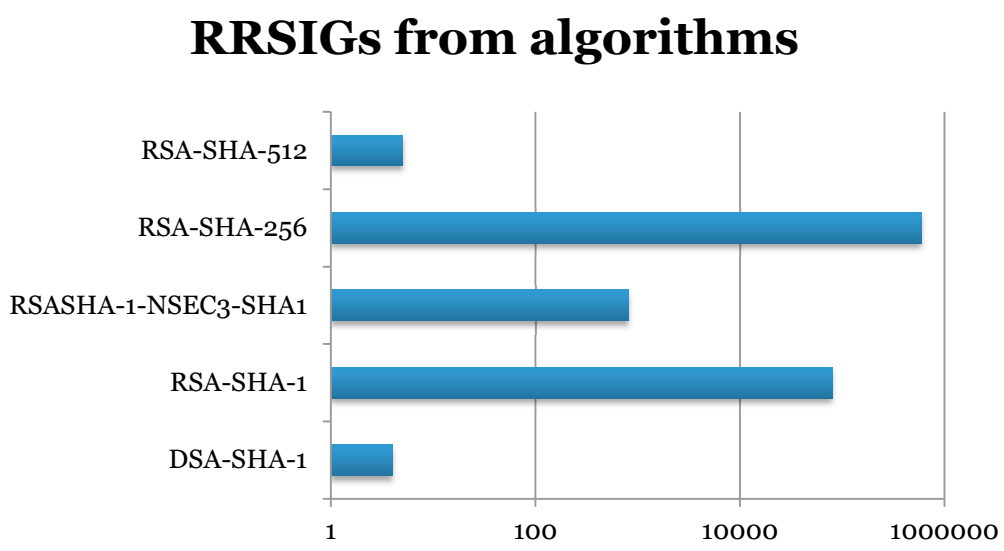

There are two types of keys in DNSSEC; Key Signing Key (KSK) and Zone Signing Key (ZSK). KSK is what you usually send to the parent zone (.se in our case), and it is published there as a fingerprint in the form of a DS record. KSK is also the key that signs the RR set DNSKEY, while ZSK signs the remaining records in the zone. The allocation between KSK and ZSK looks like this:

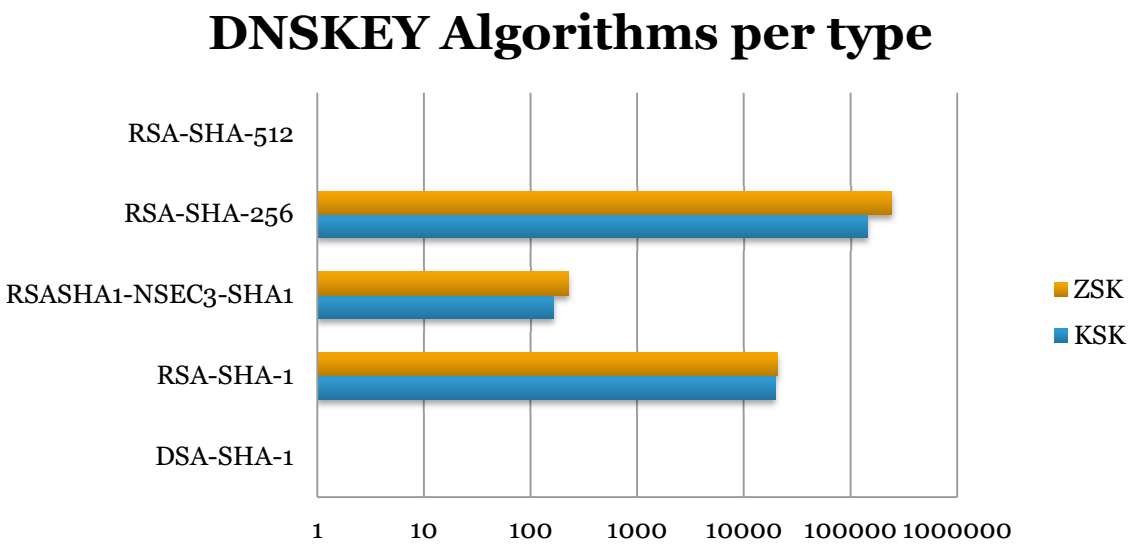

The average number of DNSKEYs per domain is 2.61, and allocated according to key type it averages 1.00 KSK per domain and 1.61 ZSK per domain. The fact that we see significantly more ZSKs per domain has to do with the fact this is the key that has traditionally been replaced most often. Since rolling a ZSK means that one must keep the old one for a while so that signatures have time to disappear from caches in the DNS resolvers, this also means that you must often use several ZSKs in parallel during a certain period of time.

The reason for the division between KSK and ZSK is that you need to integrate with the parent zone in order to switch out a key, and the longstanding belief

that switching keys often is favorable, which is up for discussion in various IETF working groups.

The practical difference between KSK and ZSK is that KSK is the one that is sent to the parent zone, and ZSK is the one that signs the remaining content in the zone. Since KSK is not replaced as often, it frequently also has a longer key length than ZSK. This is the breakdown of all of the different signed zones at .se:

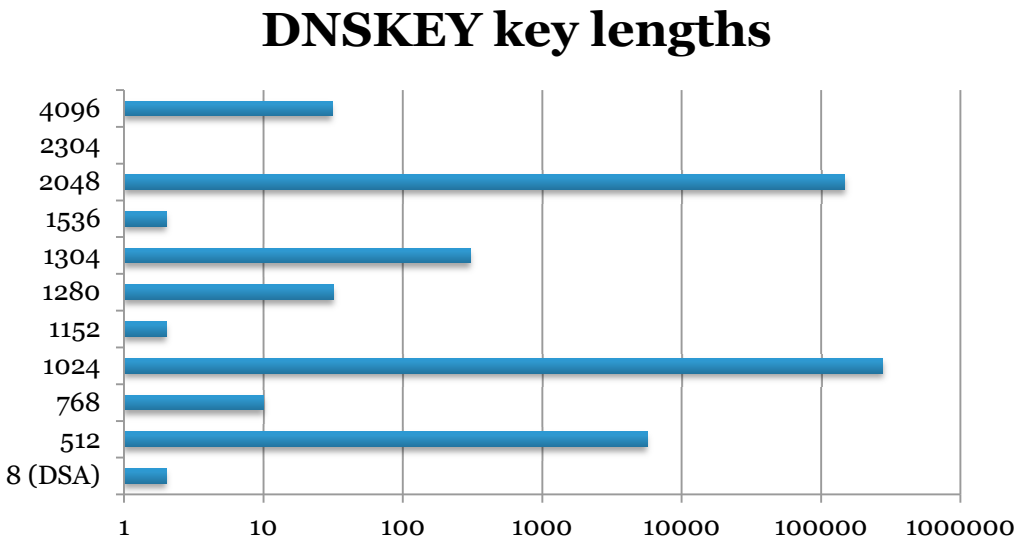

Only DSA has a key length of 8 bits, and a DSA key's key length lacks an equivalent in RSA. We will be seeing other, more unusual key lengths once ECDSA (elliptic curve cryptography) becomes standardized for DNSSEC. The graph above primarily indicates that what we see is that RSA is used for essentially all keys (since that is what is currently being used for DNSSEC).

To confirm the theory that key length for KSK is longer than for ZSK, we must break down the results for both key types:

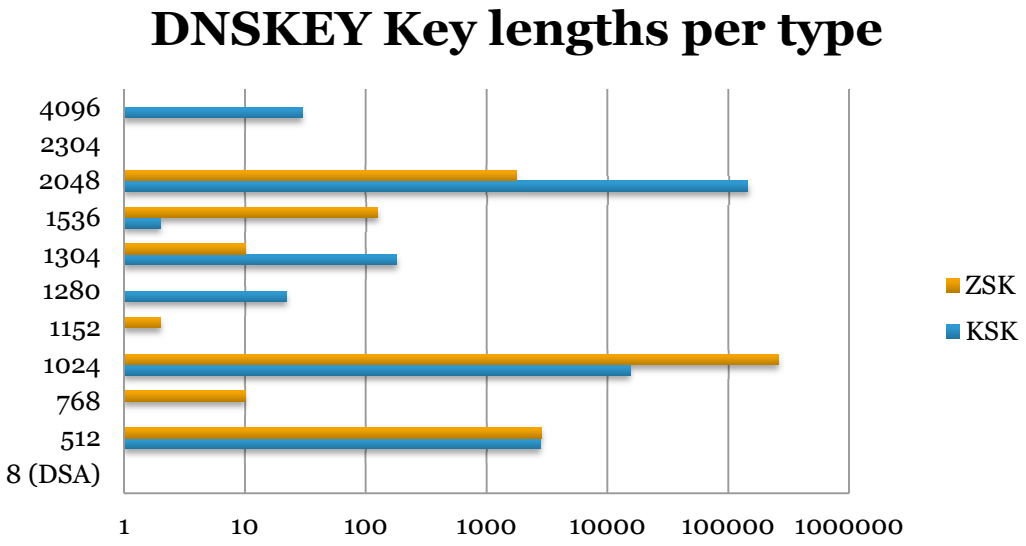

In the graph above we clearly see that KSK is far longer than ZSK.

4096 bit keys are used only by KSK, and it should be the most security-conscious (or suspicious) users that use that key length. We want to directly advise against using 512 bit keys today, which we see a relatively high prevalence of.

6.5.1 NSEC and NSEC3

NSEC and NSEC3 are the records in a DNSSEC signed zone that exist to prove that records *do not* exist. If, for example, someone tries to look up xyz.iis.se and xyz does not exist, then you also want proof of this. That is achieved by presenting signed NSEC or NSEC3 records to the resolver. These records tie together all labels (names, such as xyz) into a zone in order to build a chain that should be impossible to break.

NSEC was the first method to solve the problem with an answer to queries about domains that did not exist. As we describe above, it ties together all labels in a zone. This poses a problem, however, for those who believe that zone content is worth protecting (despite DNS being a public service), since, by using NSEC, you can relatively easily retrieve all names from a zone by "walking" along the NSEC chain, which is known as "zone walking."

One solution to this problem has been developed through NSEC3. The solution is to obfuscate names on these labels through *hashing* (calculating a checksum for) them. To also avoid having someone use "rainbow tables" (pre-generated tables with all the *hashes* of known names), you can use NSEC3 to hash a label many times ("iterations"), and even add a salt value – a randomly selected amount – to "dirty up" the label and make it particularly difficult to generate these rainbow tables in advance. The method can make it incredibly calculation intensive to try to crack a NSEC3 chain, partly due to salt length and partly due to the numerous iterations.

One problem with NSEC3, especially where you are using many iterations in the hash function, is that for an authoritative name server, it can be problematic to demonstrate a name's non-existence on a large scale. The authoritative name server must namely independently calculate the non-existence by running the hash function with both salt and the defined amount of iterations. An increased number of questions to a non-existent name will cause the server to work intensively, and ultimately it runs the risk of being subjected to a Denial of Service attack. In that instance, you must balance safety and performance, especially in the choice of the number of iterations.

NSEC3 also provides poor protection against pure dictionary attacks, since the names in a zone are often not particularly unique (meaning that they can easily be found in a dictionary). The most common domain names are so short that someone who wants to generate a list over all names in a zone can easily do so and then generate a large number of queries and in that way compile a comprehensive list of the content.

In our survey, we can see the use of NSEC in relation to NSEC3 by looking at the DNS record NSEC3PARAM. NSEC3PARAM is used to, among other tasks, hold

information about the salt data and the number of iterations, which is why we can see these values. The relationship between NSEC and NSEC3 is fairly close to 50/50, which is apparent from the graph below.

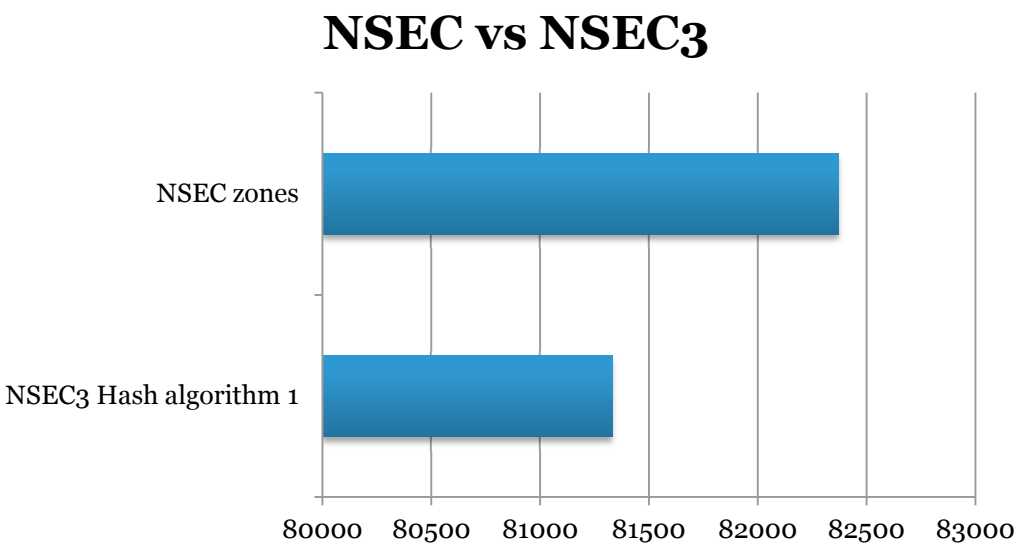

For NSEC3 zones, we are also able to look at the salt length. In terms of the content of the salt, we have not studied that in detail. A DNSSEC signer should, however, replace the salt at regular intervals to avoid an attack using rainbow tables. Since we have not looked closely at DNSSEC data over time, we are unable to report anything in this survey about how frequently the salt should be replaced.

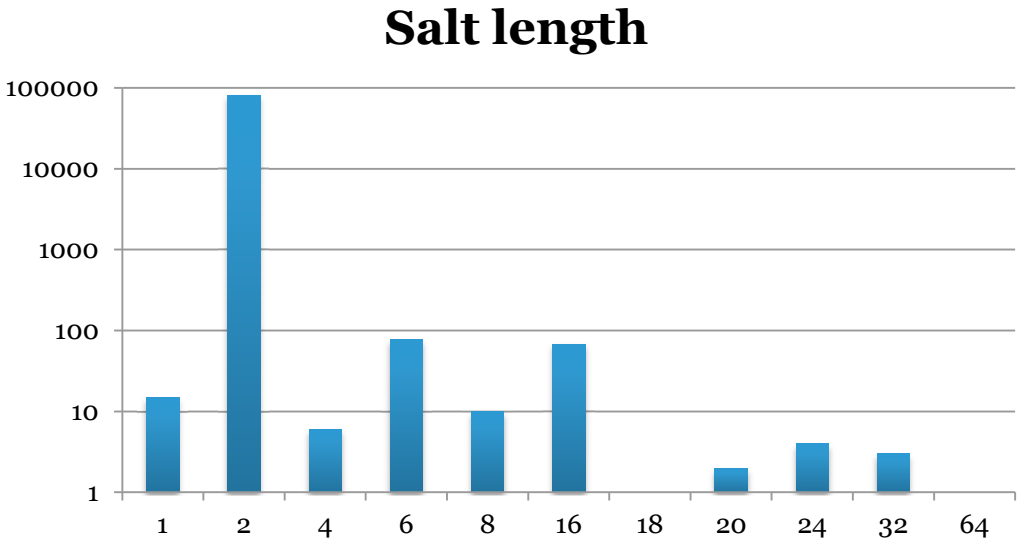

If we look more closely at the number of iterations in NSEC3, the most common amount by far is “1” (which, in reality, means that two hashes are created from the name). However, this result could be due to the fact that the same name server software handles the absolute majority of the signed domains in the .se zone. *(The software is PowerDNS, which we did not test, but rather what we*

know through strong relationships with DNS operators who have provided us with that information.)

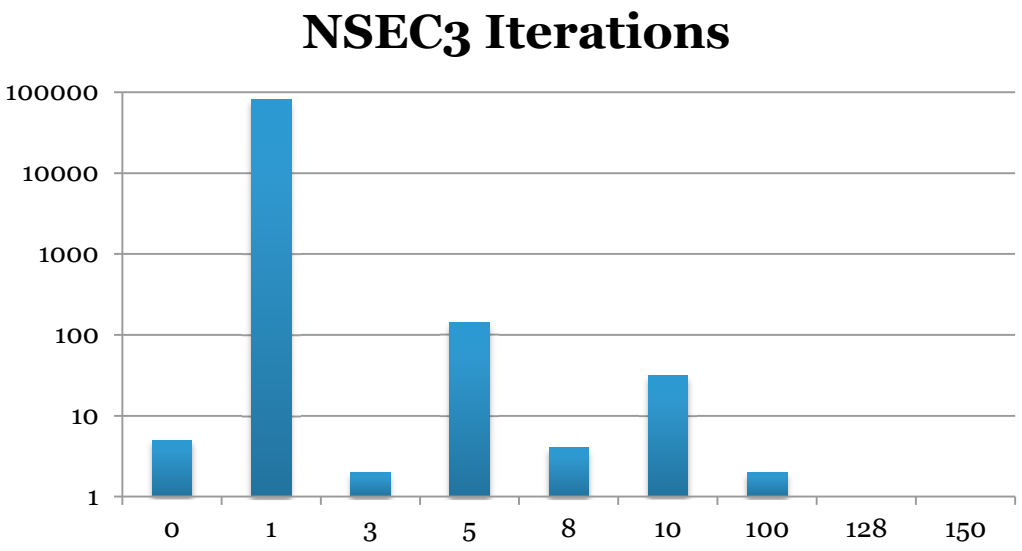

We can also conclude that there are very few domains that have an extremely high number of iterations, which negatively affects the performance of the name server during queries for responses that do not exist, making it easier to become the object of Denial of Service attacks.

6.6 DS records in the .se zone

For the entire reliance chain to work, a DS record must be published in its parent zone. These DS records are fingerprints of the zone's KSK. The fingerprints can, of course, also have several different algorithms. DS Digest type 2 is, quite simply, SHA-256, and is consequently a little newer to NDSSEC than type 1, which is SHA-1.

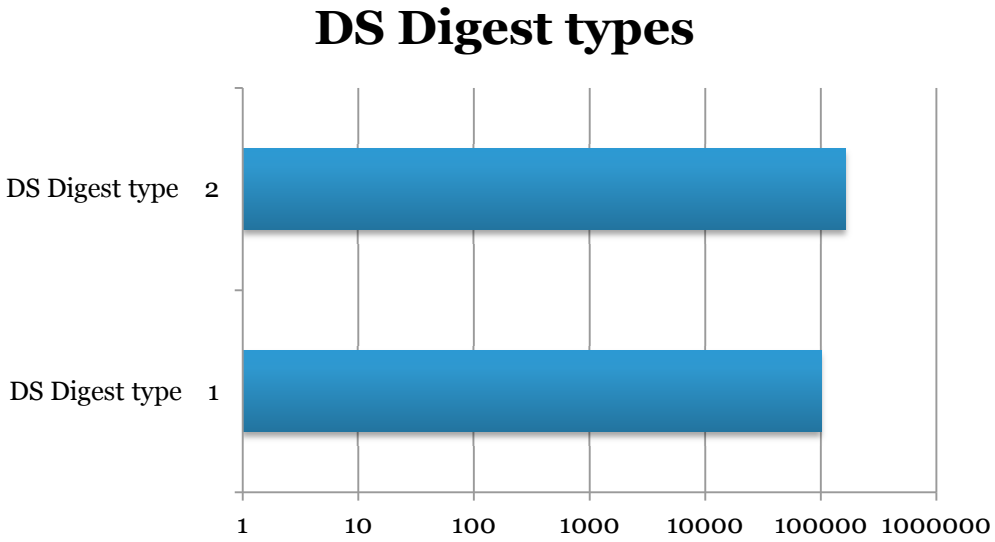

Our own registrar, .SE Direct, always publishes both types for a key in the .se zone. When .se was signed in 2005, SHA-1 was the only digest type that was

defined. Consequently, the resolvers that validated DNSSEC were incapable of verifying type 2 fingerprints. So the choice, once DNSSEC "was launched" at .se in 2007, was to always publish both DS types. Today, there are hardly any resolvers that validate DNSSEC that do not have support for the more secure type 2, thus enabling us to essentially phase out type 1. To validate with the root as a trust anchor, you must always have support for the newer algorithms, in terms of DNSKEY, NSEC3 and DS.

We also examined whether more domains share the same keys by looking at unique DNSKEY material for KSK. There are a few name server operators who use shared keys, and there are a total of 102,948 domains that share KSK with another domain, allocated as follows:

Key	Number of domains
KSK1	53,224
KSK2	43,642
KSK3	6,075
KSK4	505
KSK5	7

A manual review shows that these domains also share the same ZSKs.

The fact that one would want to share keys between several domains could be due to the fact that you have limited space for the (secure) storage of keys. If you have limited storage space for keys, or limited CPU resources to generate new keys, then shared keys can be a convenient way to solve signing for all domains since all of these keys still reside within one and the same security domain. Situations could arise where you were somehow forced to share the private key, in which case you would unfortunately affect the keys for all of the domains that share the same key.

7 Future DNSSEC surveys

In this report, we only analyzed data from a single testing occasion. Below, we outlined a number of possible examples for expanded testing that we believe would be interesting to conduct.

7.1 More frequent surveys over extended periods of time

Limiting a test to a single occasion could result in a somewhat misleading impression of signature lifespan, for example. Depending on how often a DNSSEC signer updates the signatures for the zone, the lifespan can vary widely over time, if for example resigning occurs once a week. We would be able to see this more clearly if we conducted testing once a day for a couple of weeks, for example.

If we were to conduct testing over a longer period of time than that, say once a week for six months or a year, then we would also see key rolling (especially ZSK), salt replacement for NSEC3 and similar more long-term dependent details. We would also be able to see if zones acquire temporary problems due to the fact that signatures temporarily expire, and measure the scope of the different operational problems.

7.2 New crypto algorithms

Under RFC 5933, which was published in July 2010, GOST in DNSSEC was introduced. These are Russian crypto algorithms and cover DS, DNSKEY and RRSIG. In the .se zone we do not see this algorithm in use, particularly because .SE Registry does not yet allow it. In the near future, we also expect ECDSA (elliptic curve cryptography) to be introduced in DNSSEC. The use of these algorithms could also be worth examining in further detail.

7.3 Key structure and occurrence of Combined Signing Key

The need for the splitting of keys for DNSSEC into KSK and ZSK has been increasingly debated in the DNSSEC world. The current norm is to split up the keys, although combining these into a “CSK”, Combined Signing Key actually works perfectly. We have not looked into the occurrence of CSK in this survey, but that could be worth examining. We do not believe, however, that its use is especially commonplace yet.

7.4 Quality of key material for DNSSEC

Nor have we looked at the key material for DNSSEC keys. In research results regarding SSL keys that was presented a short time ago, it was revealed that certain problems exist with prime number generation for RSA keys.[8] This means that we suspect that there could be problems with how random RSA keys are generated for DNSSEC as well. We also have historical examples of substandard random numbers. In that context, a random number bug in OpenSSL from Debian is frequently mentioned.[9]

[8] http://eprint.iacr.org/2012/064.pdf
[9] http://www.debian.org/security/2008/dsa-1571

7.5 Segmentation of the analysis

It would also be interesting to segment the analysis into smaller groups, and perhaps assimilate the annual Health status report and examine specific groups of domains, such as government agencies and municipalities in greater detail.

7.6 Summary of DNSSEC analysis

We have found domains with signature lengths that are both unexpectedly long as well as too short. NSEC3 is essentially adequate. Most domains use 2,048 bit RSA keys as KSK and 1,024 kit keys as ZSK. In the future, we anticipate that these keys may become somewhat longer. A few too many domains are using 512 bit keys, a far too insufficient choice of key length in the year 2012 in our opinion.

We can begin to discontinue the double publication of DS types 1 and 2, as the publication of type 2 is sufficient today.

All too often, SOA Expire lacks a connection to RRSIG expiration time, these parameters should definitely be reviewed.

Appendix 1 - About .SE and the survey

According to its charter, the purpose of .SE (the Internet Infrastructure Foundation) shall be “to promote positive stability in the Internet infrastructure in Sweden and to promote research, training and education in data and telecommunication, with a specific focus on the Internet. By so doing, the Foundation must assign priority areas that increase the efficiency of the infrastructure for electronic data communication, whereby the Foundation shall, inter alia, disseminate information concerning R&D efforts, initiate and implement R&D projects and implement high-quality inquiries.” Secure Internet infrastructure is a very important and key area for us.

The considerable interest shown in the results of the studies of earlier years convinces us at .SE that the study is valuable and we will continue to conduct it, this year going more in-depth in certain areas. It is part of a long-term project called the Health Status Project.

.SE has been responsible for the operation and administration of all name servers for .se domains since 1997 and, over the years, has amassed solid experience with regard to the domain name system (DNS). International Best Common Practice for DNS has gradually emerged from the organization’s mistakes and experiences, and those of other parties and this practice can also be applied to environments other than only top-level domains. DNS is somewhat of an unknown system that has existed for nearly 30 years. Throughout the years, DNS has proven to offer exceptional scalability and robust design. Essentially no changes have been required in the basic protocols, despite the enormous growth of the Internet. However, DNS has become increasingly important to the existence of functioning communication between Internet users worldwide, and this requires that all areas of DNS maintain a high level of quality.

Appendix 2 - Industry standard for high-quality DNS service

For the more technically skilled reader, we have provided a more detailed description of the industry standard for high-quality DNS service in terms of recommendations in this appendix. You can easily test your domain yourself on .SE's website.

Our DNSCheck tool can also perform what are known as undelegated domain tests. An undelegated domain test is a test carried out on a domain that can be (but does not have to be) published entirely in DNS. This function is highly useful for those who want, for example, to relocate a domain from one name server operator to another. For instance, let us say that the domain example.se is to be relocated from the name server "ns.nic.se" to the name server "ns.iis.se". In this case, an undelegated domain test can be carried out on the domain (example.se) using the name server to which the domain will be moved (ns.iis.se) BEFORE the move itself is implemented. When the test shows a green light, it is relatively certain that the domain's new home at least knows that it should respond to queries regarding the domain. However, errors in the zone information may still exist and may not be detected by this test.

This function is available in both Swedish and English at:

http://dnscheck.iis.se/

1. At least two name servers

Recommendation: DNS data for a zone should be located on at least two separate name servers. For reasons of availability, these name servers should be logically and physically separated so that they are located in different service-provider networks in different autonomous systems (AS).

Explanation: At least two functioning name servers should exist for each underlying domain. They should be listed as NS records for the domain in question. They should be physically separated and located in different network segments to obtain optimum functionality. This will ensure that the domains continue to function even if one of the name servers stops working.

Consequence: When the sole server or sole service provider experiences a disruption, DNS service will be rendered unreachable for the domain on that server or in the service provider's network. Accordingly, the services under the domain will not be reachable, even if they are located with entities other than the organization's own name server operator.

2. All name servers specified in a delegation should exist in the underlying zone

Recommendation: All of the NS records listed in the overlying (parent) zone (.se or equivalent) in order to point out (delegate) a certain domain should also simultaneously exist in the underlying (child) zone.

Explanation: NS records are used in the overlying zone to transfer responsibility for (delegate) a certain domain to other servers. According to DNS documentation, this list of computers should also be found in the zone file that "receives" the responsibility and that contains other data about the zone. The lists must be kept synchronized so that all NS records included in the

parent zone are also found in the child zone. The list in the parent zone is not automatically updated; it is only updated after a “manual” report is submitted to the responsible registration unit. If changes are required that entail a change to the overlying zone, the administrative contact for the underlying zone shall immediately inform the registration unit.

Consequence: If the parent zone contains information about the child zone that de facto does not exist in the child zone, this means that anyone submitting queries about the domain will not receive a response, thus resulting in an impact on availability.

3. Authority

Recommendation: All name servers listed with NS records in a delegated zone shall assume authoritative responsibility for the domain.

Explanation: When checking the sub-domain servers, it should be possible to obtain consistent and repeatable authoritative responses for SOA and NS records for the sub-domain. This applies to all servers listed in the underlying zone’s DNS for the domain in question.

Consequence: DNS usually functions even if this defect exists. However, a defect existing in a zone indicates weaknesses in the procedures of the party responsible for the content of the domain’s DNS.

4. Serial numbers for zone files

Recommendation: All name servers listed with NS records in the delegated zone shall respond with the same serial number in the SOA record for the domain.

Explanation: The serial number in the SOA record is a type of version number for the zone, and if the servers have the same serial numbers for their zones, this indicates that they are synchronized. This is controlled by sending SOA-record queries to each server and comparing the serial numbers of the responses. SOA is the acronym for Start of Authority.

Consequence: If the name servers are not synchronized and do not have the same version of the zone file, the entity submitting a query about a domain risks not receiving a response. Availability will be affected.

5. Contact address

Recommendation: The zone contact address in the SOA record must be reachable.

Explanation: The SOA record for a domain includes, along with other sub-records, an e-mail address that is to serve as a contact point if the administrator of the domain in question needs to be reached. In simple checks, e-mail servers for the e-mail address shall not provide obvious error messages (for example “user unknown”). In more detailed checks, it should be possible to send test messages to the address and receive responses to these within three days.

Consequence: The reason for having a current e-mail address for contacts is that it must be possible to quickly call attention to problems relating to the

reachability of a domain. If such an address does not exist, it will become more difficult to solve problems arising in DNS due to an individual domain.

6. Reachability

Recommendation: All NS records in the underlying zone must be reachable for DNS traffic from the Internet.

Explanation: The NS records for a domain comprise the list of the computers that function as name servers for the domain. All listed servers must be reachable via the Internet at all of the addresses listed in the corresponding address entries in DNS for the computers in question.

Consequence: If a name server is not reachable despite its name being included in the list of name servers that respond to queries about a domain, this means that entities submitting queries will not receive responses. Availability will be affected.

7. SOA parameters

Recommendation: Follow RIPE's recommendations regarding parameters in the SOA record:

```
example.com.  3600  SOA  dns.example.com. hostmaster.example.com. (
          1999022301   ; serial YYYYMMDDnn
          86400        ; refresh (  24 hours)
          7200         ; retry   (   2 hours)
          3600000      ; expire  (1000 hours)
          172800 )     ; minimum (   2 days)
```

Explanation: These parameters control how the responsible name servers handle the zone, and are explained in greater detail in Section 4 of RIPE-203:

http://www.ripe.net/ripe/docs/ripe-203

However, new for DNSSEC is that expire should be longer (approximately 2/3) than the shortest RRSIG Expiration.

Consequence: In the event of network problems, for example, the zone could become involuntarily disconnected.

Appendix 3 – Information on DNSSEC

When DNS was created in the 1980s, the main idea was to minimize central administration of the network and make it easy to connect new computers to the Internet. However, no major importance was attributed to security. The deficiencies in this area opened the way for various types of abuse and attacks where the responses to DNS lookups are falsified. This way, Internet users can be misguided; for example, people can be tricked into disclosing sensitive information such as passwords and credit card numbers. Some of the most well-known and greatest threats to DNS are cache poisoning and pharming.

Cache poisoning is a situation whereby, either by attack or inadvertently, DNS data is introduced into a name server that did not originate from an authoritative source. One of the most notorious examples of this was the much discussed Kaminsky bug in 2008.

Pharming is a when someone makes the actual DNS content point to the wrong servers. This specifically means that an Internet address for a bank, for example, may be redirected to an entirely different server, although for the visitor, the address field still makes it appear as though he/she is visiting the right server. Accordingly, there is no doubt that DNSs need to become more secure. DNSSEC is a long-term solution that protects against several different types of manipulation of DNS queries and responses transmitted between different servers in the domain name system.

To counteract these types of attacks, security extensions have been developed for DNS that are designated DNSSEC (DNS Security Extensions). DNSSEC is based on cryptographic keys that are used to sign the content of the zone files. The validation of signatures ensures that the responses truly derive from the right source and have not been changed during transmission. The increased security that DNSSEC provides means that many attacks no longer have any effect.

.SE's launch of DNSSEC service for more secure DNS in 2005 has also contributed to a greater focus on DNS and DNS operation. Companies wishing to make their DNS infrastructure more secure by using DNSSEC realize relatively quickly that they cannot introduce the mechanism until they first review their own DNS infrastructure as a whole.

Accordingly, we are naturally interested in finding out how well prepared .se domains are for DNSSEC. This – as well as the fact that we are responsible for the Swedish top-level domain – is the crucial reason why our tests focus specifically on the quality of DNS.

The signing of what is known as the root zone in summer 2010 accelerated the proliferation of DNSSEC. The root zone's location at the pinnacle of DNS hierarchy facilitates the implementation of DNSSEC for the underlying top-level domains.

One of .SE's employees was selected as a Trusted Community Representative (TCR) in order to act as a Crypto Officer (CO) and participate in the key ceremonies that are performed for the root zone four times a year; twice at the site located on the west coast of the US and twice at a corresponding site on the east coast of the US.

Unlike the traditional domain name system (DNS), DNSSEC look-ups have a cryptographic signature, which makes it possible to ensure that these look-ups come from the right user and that the content is not changed during transmission. The aim of the service is to ensure that domain registrants can secure their domains using DNSSEC.

What DNSSEC protects against

The purpose of DNSSEC is to safeguard the content of DNS using cryptographic methods requiring electronic signatures. Through the validation of signatures, DNSSEC allows the user to determine whether the information returned from a look-up in DNS comes from the correct source and whether it has been manipulated en route. Thus, it is difficult to falsify information in a DNS that is signed with DNSSEC without it being detected.

For ordinary users, DNSSEC reduces the risk of being defrauded, for example, when conducting bank transactions or shopping on the Internet, since it is easier for the user to determine whether he or she is really connected to the correct bank or store and not to an impostor.

However, it is important to note that DNSSEC does not stop all types of fraudulent activity. It is only designed to prevent attacks in which attackers manipulate responses to DNS queries for their own gain.

What DNSSEC does not protect against

A number of other security issues and problems on the Internet remain that DNSSEC cannot solve, including Distributed Denial of Service (DDOS) attacks.

DNSSEC provides some protection against phishing (websites that resemble or are identical to genuine websites to trick users into revealing passwords and personal data), pharming (redirecting a DNS query to the wrong computer) and other similar attacks against DNS. DNSSEC does not prevent attacks at other levels, such as at the IP or network level.

.SE's role in DNSSEC

Many have been waiting for the root zone, meaning the parent zone of .se, to be signed and this became a reality in 2010. To date, .SE has been responsible for signing .SE's zone file and for acting as a *trust anchor* in the chain for the

Swedish part of the Internet. A *trust anchor* signs the keys of the underlying zones and acts as the starting point in the verification chain. Signing means that .SE assumes responsibility for managing and verifying the DS records of the underlying zones. This is comparable with the management of NS records in DNS.

.SE will still sign .SE's zone file, although since .SE publishes its DNSSEC keys in the root zone, it is now the root that constitutes the *trust anchor* for the Internet. This makes it easier for all resolver operators that would otherwise be forced to manage all keys for all signed top domains, which are *trust anchors* for each of their underlying domains. With the root signed, they only need to keep track of the root key. Modern standards also offer simpler management of key exchanges and new tools has been developed to make it easier (refer to Open DNSSEC below).

Further information on .SE's DNSSEC service is available at https://www.iis.se/en/domaner/dnssec/

.SE provides additional information on DNS vulnerabilities at https://www.iis.se/en/domaner/dnssec/kaminskybuggenThe website's features include a link to a film that demonstrates how an attack is carried out and the ability to test whether the resolver being used is vulnerable to the Kaminsky bug.

Here are some links to further information:

Information on DNSSEC and the advances in both its use and tools.
http://dnssec.net

A practical guide on how to implement DNSSEC.
http://www.nlnetlabs.nl/publications/dnssec_howto/index.html

News from DNSSEC Deployment Initiative is distributed regularly at
http://www.dnssec-deployment.org/

The Initiative also has an e-mail list that anyone can subscribe to and thus stay abreast of developments in the field.

OpenDNSSEC

DNS is relatively complex, as are electronic signatures. Naturally, the combination of these in DNSSEC is also complex.

After .SE noted that the lack of high-quality, accessible tools in the market for signing zone files with DNSSEC was a barrier for many parties who wished to start implementing DNSSEC, a development project was launched in conjunction with some of the foremost developers in the area. The result was OpenDNSSEC, which is a turnkey program, or a tool for facilitating the implementation and use of DNSSEC. OpenDNSSEC secures DNS information the moment before it is published on an authoritative name server. OpenDNSSEC takes an unsigned zone file, adds signatures and other items for DNSSEC and sends the file on to the authoritative name servers for the relevant zone.

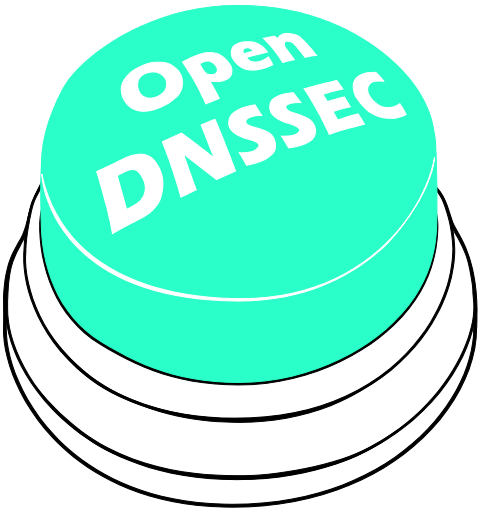

The purpose of OpenDNSSEC is to manage these difficulties and relieve system operators of responsibility for them once the operators have set up the system.

By participating in the development of a turnkey system for signing zone files with DNSSEC, .SE hopes to facilitate the spread of DNSSEC.

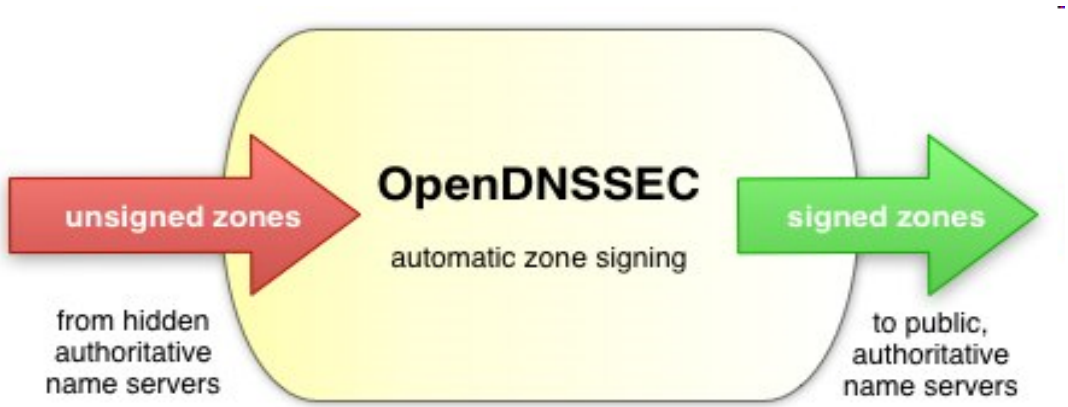

OpenDNSSEC was developed under a special company owned by .SE (The Internet Infrastructure Foundation).

OpenDNSSEC is the result of collaboration between developers from .SE, Nominet, NLNet Labs, SIDN, SURFnet, Kirei AB and Sinodun. More information is available at http://www.opendnssec.org/

The software, which is openly available, can also be downloaded and tested from the website.

.SE (The Internet Infrastructure Foundation) is a not-for-profit public-service organization that acts to promote the positive development of the Internet in Sweden. .SE is responsible for the Internet's Swedish top-level domain, .se, encompassing domain-name registration and administration, as well as the technical operation of the national domain name registry. Proceeds from domain-name registrations are used to support projects that contribute to the Internet development in Sweden, through proprietary operations and the financing of independent projects.

This survey is included in of .SE's Health status focus area. The aim of this focus area is, among other things:

- To monitor the quality of the Internet's infrastructure in Sweden by compiling and analyzing facts,
- To disseminate the results from the surveys, and
- To use advice and recommendations to contribute to ensuring that the infrastructure functions well and has a high level of accessibility.

Another aim is to, when necessary, detect deficiencies and improprieties.

.SE (The Internet Infrastructure Foundation)
P.O. Box 7399, SE-103 91 Stockholm, Sweden
Tel +46 (0)8 452 35 00, Fax +46 (0)8 452 35 02
Org. nr 802405-0190, www.iis.se, info@iis.se

.se
Moving the Internet forward