UNIVERSITY OF AMSTERDAM

MASTER THESIS

SECURITY AND NETWORK ENGINEERING - RESEARCH PROJECT II

Verifying email security techniques for Dutch organizations

July 11, 2018

KEYWORDS: EMAIL, SECURITY, SPF, DKIM, DMARC, STARTTLS, DNSSEC, DANE

Vincent van Dongen
vincent.vandongen@os3.nl

Supervisors
NLnet Labs - Ralph Dolmans
NLnet Labs - George Thessalonikefs

Abstract

Email security hasn't been taken into account during the original design of email protocols. Therefore, different techniques have emerged to secure email communication and to validate emails. Many governments have defined guidelines that require or strongly recommend to adopt these techniques to improve email security. This research investigates which and how many of these techniques have been adopted by organizations within the Netherlands. A list of Dutch organizations with more than 10 employees has been scraped from two websites (in total 46.650 unique organizations) that was used as an input for the experiment. The Dutch organization named Standardization Forum has defined a list of compulsory email security techniques. A tool from 'internet.nl' was used to check whether or not these email security techniques have been adopted by Dutch organizations.

We found a relation between the type of sectors. The 'Public service' sector has the highest score. We assume that the high score is related to compulsory policies for governmental organizations because many governmental organizations are present in the 'Public services' sector.

Contents

Acknowledgement

I would like to thanks Ralph Dolmans and George Thessalonikefs from NLnet Labs for supervising this research project.

1 Introduction

Email has become an important tool for organizations to communicate and exchange (sensitive) information. Email security hasn't been taken into account during the original design of email protocols. Therefore, different techniques have emerged to secure email communication and to validate emails. Many governments have defined guidelines that require or strongly recommend to adopt these techniques to improve email security [14] [27] [31]. Some studies have shown that not every organization or mail provider has adopted these techniques [17]. This research investigates which and how many of these techniques have been adopted by organizations within the Netherlands.

For this research, a list of Dutch organizations was created that was used to verify whether or not email security techniques have been adopted by Dutch organizations. Additional information about the organization such as the number of employees and the location was also added to the list. This list was created by scraping and parsing information about the organization from the internet. The website 'internet.nl' has built a tool that already verifies many email security techniques for a given domain [19]. This tool was used during the experiment and the created list of Dutch organizations was used as input for this tool. The results from the experiment were analyzed to determine how many email security techniques have been adopted by organizations within the Netherlands.

1.1 Research question

The main research question that was defined for this project is as follows:
How many email security techniques have been adopted by organizations within the Netherlands?

The following sub-questions have been defined to answer the research question:

1. Which techniques do exist to secure email?
2. What is the most feasible way to create a data-set of Dutch organizations that also contains the number of employees and the type of sector per organization?

   The following sub-questions will be answered based on the results:

3. Is there a distinction between the different sectors regarding the adoption of email security techniques?
4. Is there a distinction between the size of an organization regarding the adoption of email security techniques?
5. Is there a geographical distinction between organizations regarding the adoption of email security techniques?

1.2 Related work

Previous research has been done on verifying SPF, DKIM and DMARC records [18] [6]. The results from the reports show that not every mail provider has adopted SPF, DKIM and DMARC. The research didn't specify which type of organizations were verified. In 2015, the Dutch Internet Standards Platform launched a tool on the website (internet.nl) to check if an email server is compliant [19]. This tool was used during the experiment.

1.3 Scope

For this research, only Dutch organization will be verified [3]. The website 'internet.nl' already checks for email security. Therefore, a new tool does not need to be built.

1.4 Approach

This research is divided into different parts. The first part of this report discusses how email works, which techniques exist to secure email and how you can verify whether or not these techniques have been adopted. Next, a data-set of Dutch organizations has to be created. This part also discusses what type of data should be present in the data-set to answer the research questions. Next, an experiment is conducted that uses the data-set as input for the tool. Finally, the results that were generated during the experiment are discussed and the research questions are answered.

2 Background information

This chapter briefly discusses how email works, which techniques exist to secure email and which of these techniques will be verified for this research.

2.1 Theoretical background

A number of components are used to create, send and transfer emails. A program which allows someone to send and receive e-mails is known as a Mail User Agent or MUA. A MUA is a software component (or web interface) that allows an end user to compose and send messages to one or more recipients. Examples of MUAs are Mozilla Thunderbird and Microsoft Outlook. When a MUA sends a message, it uses a Mail Submission Agent (MSA). The MSA receives the email messages from a MUA and cooperates with a mail transfer agent (MTA) for delivery of the mail. The MTA then checks the message to determine the recipient and queries the Domain Name System (DNS) servers to find out which other MTA is responsible for handling e-mail for the recipient. It then sends the message to that MTA. An email message may pass through multiple MTAs before reaching the final destination. Eventually, the MTA will pass the message to a Mail Delivery Agent (MDA). The MDA is responsible for actually storing the message to disk [8] [28] [24].

2.2 Techniques to secure email

Different techniques exist to secure email. The Dutch Standardization Forum (part of the Minister of Economic Affairs) has defined a list of compulsory standards. The list of mandatory standards also contains different email-related security standards that all Dutch governmental organizations must adopt and implement in their email infrastructure. [24]. The following email-related security standards that are listed in the mandatory standards will be verified during the experiment. Each technique will be discussed in the next paragraphs.[1]

Techniques	Check	Checks if
SPF	Record available	An SPF record is available
	Policy	A sufficiently strict policy is used
DKIM	Record available	A DKIM record is present
DMARC	Record available	A DMARC record is available
	Policy	A sufficiently strict policy is used
DNSSSEC	Signed domain	The domain is DNSSEC signed
	Secure domain	The domain is signed with a valid signature
	Signed mx record	The domains the MX records point to are signed
	Validate signed mx record	The domains the MX records point to are signed
DANE	Record available	Each of the mail server domains provide a TLSA record for DANE
	Valid record	The DANE fingerprint presented by the mail server domains are valid for the mail server certificates.
STARTTLS	Supports	The mail server supports the STARTTLS option
	TLS version	The mail server supports sufficiently secure TLS versions
	Cipher suites	The mail server supports sufficiently secure cipher suites
	Trust chain of certificate	A valid chain of trust can be build from the certificate
	TLS compression	The mail server supports TLS compression
	Public key of certificate	The bit-length of the public key of the mail server certificate is sufficiently secure
	Signature of certificate	The signed fingerprint of the mail server certificate was created with a secure hashing algorithm
	Domain name on certificate	The domain name of the receiving mail server matches the domain name of the certificate

Table 1: A list of standards that will be verified during the experiment.

[1] list of mandatory standards defined by the Dutch Standardization Forum: https://www.forumstandaardisatie.nl/open-standaarden/lijst/verplicht

2.2.1 Sender Policy Framework (SPF)

SPF was designed to address emails being sent by unauthorized senders. When an MTA receives an email from another MTA, the receiving MTA can check the IP-address of the connecting MTA. The receiving MTA then checks the domain part of the envelope From-address and queries the related DNS server to request the SPF record. An SPF record contains the IP-addresses of MTAs that are authorized to send email from that domain. The IP-address from the sending MTA must match one of the IP-addresses in the SPF record. This technique ensures that emails are only sent from authorized domains that are listed in the SPF record. An SPF record can contain different mechanisms to identify a set of IP-addresses that are permitted or not permitted for sending mail. Note, that an SPF record can include more mechanisms than only IP addresses [24]. An SPF record contains a parameter called qualifiers in which a policy is defined for a mechanism. When the mechanism 'ALL' is defined, the qualifiers should be either - or due to the fact that the qualifiers + and ? in combination with the mechanism 'ALL' actually allow every IP-address that hasn't been matched with the previous mechanism. Therefore, the mechanisms 'ALL' in combination with '+' and '?' can be considered as insufficient [20] [16].

Two examples of SPF records are displayed below. The first example only allows the sending MTA to have the IP-address in address block 192.0.2.0/24. If the sending MTA doesn't have an IP-address in the address block, the mail shouldn't be accepted. In the second example, the sending MTA must match the listed hosts in the MX record. However, if the sending host doesn't match, the mail is still accepted due to the ? qualifier. This can be considered as insufficient because email can still be received from unauthorized senders.

```
example.org IN TXT "v=spf1 ip4:192.0.2.0/24 -all"

example.org IN TXT "v=spf1 mx ?all"
```

Figure 1: Two SPF examples.

Parameters that are being verified:

During the experiment, two parameters regarding SPF will be checked. Both parameters can be retrieved by querying the DNS server.

- **SPF record available:** To check if an SPF record is available. The test will result in a failure if no valid SPF record is available on the DNS server. Note, by having more than one SPF record in the same domain is not valid and will also lead to a test failure [16].
- **SPF policy:** To check if a sufficient policy is used. A sufficient strict policy has to contain either all (softfail) or -all (hardfail). The test also follows the include mechanism and redirect modifier to determine if the SPF record is sufficient. If the include or redirect domain consists of macros, they are not followed as the tool doesn't have the necessary information from an actual mail or mail server connection to expand those macros.

2.2.2 DomainKeys Identified Mail (DKIM)

DKIM is a standard that ensures that the message hasn't been altered between the sending and receiving MTA servers. Furthermore, DKIM also enables the owner of a domain to claim that the message originated from an MTA that is related to its domain. DKIM uses public-key cryptography to sign emails with a private key as it leaves a sending MTA. The receiving MTA can then use a public key that was published by its DNS server to verify the source of the message, and that the body of the message hasn't changed during transit. Once the signature (created with the private key) is verified with the public key by the receiving MTA, the message passes DKIM validation and is considered authentic [29].

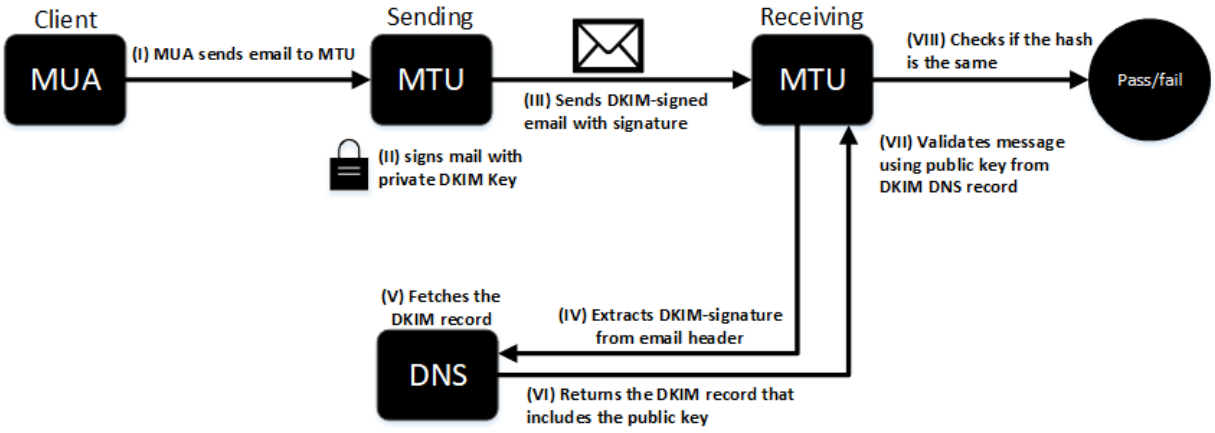

Figure 2: An overview of DKIM.

A DKIM signature is generated by the sending MTA using the email message body in combination with headers and places the signature in the header of the message along with necessary information for the client to validate the signature. An example of a DKIM signature in an email is displayed below. The tag 'v' indicates the version of DKIM, which is currently 1. The tag 'a' defines the algorithm that is used, which has to be either RSA-SHA1 or RSA-SHA256. The tags 'd', 'i' (optional) and 's' (also known as a selector) are used to form the DNS query that obtains the key that can validate the DKIM signature. The tag 't' is the time the DKIM signature was generated and the tag 'b' is the actual signature encoded in base64 format [5].

```
DKIM-Signature: v=1; a=rsa-sha256; d=example.org; i=sender.example.org;
t=1425066094; s=adminkey; b=<base64 string>
```

Figure 3: An example of a DKIM signature in a email.

When the receiving MTA gets the message, the receiving MTA attempts to validate the signature by looking for the public key indicated in the DKIM signature. The receiving MTA issues a DNS query for the DKIM record. An example of a DKIM record is displayed below. The 'v' tag in a DKIM record is also used to define the DKIM version. The tag 'k' defines the key type. RSA is currently the only specified algorithm that is used for the key type. The tag 'p' is the actual public key encoded in base64 format [15].

```
adkimkey._domainkey.example.org IN TXT "v=DKIM1; k=rsa; p=<base64 string>"
```

Figure 4: An example of a DKIM record in a DNS server.

Parameters that are being verified:

To verify whether DKIM has correctly been implemented, you need the DKIM record from the DNS server along with the DKIM signature that is located in the email header. In contrast to SPF, you must receive an email from the sending MTU that signs the mail with the private DKIM key to completely verify DKIM. There is no feasible solution to query and evaluate the public key in your DKIM record, because the DKIM selector (that should be in the emails you send) is needed to do so. Therefore, only the DNS server will be queried to check if a domainkey sub-record is available. This means that the DNS server must respond with a NOERROR to a query for _domainkey.<domainname>. Note, that it's possible to publish DKIM records under a different subdomain or a third party domain, but will result in a fail during the experiment. This means that it might happen that DKIM has correctly been implemented, but will fail the test.

2.2.3 Domain-based Message Authentication, Reporting and Conformance (DMARC)

SPF and DKIM were designed to provide domain-level authentication. However, neither SPF or DKIM include a mechanism to tell receivers if SPF or DKIM is in use. Furthermore, SPF and DKIM don't have a feedback mechanism to inform domain owners about the effectiveness of their authentication techniques. DMARC allows email sending domain owners to specify policies on how receivers can verify the authenticity of their email and how the receiver can handle emails that fail to verify. Furthermore, DMARC also provides a mechanism that allows receivers to send reports to the domain owner. These reports contain the sending source along with information about the message that has passed or failed SPF and DKIM.

To deploy DMARC, the sending domain owner will publish SPF and/or DKIM records in the DNS server. The domain owner also publishes a DMARC policy in the DNS server. Similar to SPF and DKIM, the DMARC policy is defined by tags. [7] [24] [9].Two important tags for DMARC are RUA and RUF. RUA addresses to which aggregate feedback is to be sent and RUF addresses which message-specific failure information is to be reported. Both addresses should be valid email addresses.[7] Two examples are displayed below. The first example contains a policy that recommends no specific actions. Moreover, the policy also doesn't request the receiver to send any feedback. Therefore, this policy can be considered as ineffective or insufficient. The second example contains the policy to reject emails that fail the SPF/DKIM validation and also wants to receive aggregate reports [29].

```
_dmarc.example.org 3600 IN TXT "v=DMARC1; p=none;"

_dmarc.example.org 3600 IN TXT "v=DMARC1; p=reject; RUA=report@example.org;"
```

Figure 5: Two DMARC examples.

Parameters that are being verified:

In order to pass the DMARC test, at least one of the two available validation methods (SPF or DKIM) should have been correctly implemented. During the experiment, two parameters regarding DMARC will be checked. Both parameters can be retrieved by querying the DNS server.

- **DMARC record available:** To check if the DNS server contains a valid DMARC record. The test will result in a failure if no valid DMARC record is available on the DNS server. Having more than one DMARC record in the same domain is not valid and will lead to a failure [24].
- **DMARC policy:** To check if a strict policy is used. A strict policy has to contain either 'p=quarantine' or 'p=reject'. The policy (p=none) is considered insufficient and will fail the test. The experiment also checks whether the mail addresses under rua= and ruf= are valid and authorized to receive DMARC reports.

2.2.4 Domain Name System Security Extensions (DNSSEC)

DNSSEC protects users from forged DNS data by using public key cryptography to digitally sign authoritative zone data. DNS groups all the records with the same type, class or name into a resource record set (RRset). This RRset could get digitally signed by the private part of the 'Zone-Signing Key pair' (ZSK) and stores them in their name server as RRSIG records. The public part of the ZSK has been made available by publishing the public key in a DNSKEY record. Another key pair is defined to facilitate key rollovers and DS record geneation and communication with the parent domain: 'Key-Signing Key pair' (KSK). The private part of the KSK signs the KSK DNSKEY and also stores them in an RRSIG record. The public part of the KSK has also been made available by publishing the public key in a DNSKEY record.

A zone operator hashes the DNSKEY record containing the public KSK and gives it to the parent zone to publish the hash as a DS record. To check the validity of the child zone's public KSK, the resolver hashes the public KSK key and compares it to the DS record from the parent. If they match, the resolver can assume that the public KSK hasn't been tampered with. The DS record is signed just like any other RRset. It's digitally signed by the private part of the ZSK and stores them in their name server as RRSIG records. This 'chain of trust' relationship between the child and the parent is repeated up to the root of DNS. DNSSEC signatures are validated by following this chain of signatures to a ''trust anchor''. A trust anchor is a DNSKEY (usually a KSK) that is placed into a validating resolver, which enables the validator to cryptographically validate the results for a given request back to a known public key. [4] [32] [25]

Email related DNS-records such as MX-records, SPF-record, DKIM-record or a DMARC-record might be present in a DNS server. In order to ensure that these records have not been altered during transmission, DNSSEC has to be implemented. Therefore, DNSSEC plays an import role in securing email and will be verified during the experiment.

Parameters that are being verified:

During the experiment, four parameters regarding DNSSEC will be checked. Both parameters can be retrieved by querying the DNS server.

1. **Signed domain:** To check if the domain is DNSSEC signed.
2. **Secure domain:** To check if the domain is signed with a valid signature.
3. **signed MX-record:** To check if the domains the MX-records point to are DNSSEC signed.
4. **validate signed MX-records:** To check if the domains the MX-records point to are signed with a valid signature

2.2.5 STARTTLS for SMTP

By default, SMTP servers and clients communicate in plain text over the internet. TLS has been introduced to encrypt traffic. STARTTLS is a way to take an existing insecure connection and upgrade it to a secure connection using TLS. Therefore, the server can offer SMTP services over plain text and an SMTP service over TLS on a single port, rather than requiring separate port numbers for secure and plaintext operations.

To use STARTTLS in an SMTP session, the SMTP sender establishes a TCP connection to the SMTP receiver. After the TCP session has established, the SMTP sender sends the STARTTLS command to the SMTP receiver. Both parties must immediately decide whether or not to switch to TLS mode [26]. The client must then send a TLS handshake to the server. After the TLS handshake has completed, the connection between both parties is encrypted. During the TLS handshake, different options such as TLS version, algorithms, cipher suites and key length are negotiated. If the options that are used during the TLS negotiation are deemed to be not strong enough or if the authentication is not good enough for either party, the client or server may choose to end the SMTP session [11].

Parameters that are being verified:

The Dutch National Cyber Security Center (NCSC) has defined guidelines which options should and shouldn't be used for a secure TLS session[2]. During the experiment, the tool from 'internet.nl' will check whether sufficient options are used that are inline with the NCSC guidelines [23]:

1. **STARTTLS available:** Checks if the receiving mail server supports the STARTTLS option.
2. **TLS version:** Checks if the receiving mail server supports sufficiently a secure TLS version.
3. **Cipher suites:** Checks if the mail server supports sufficiently secure cipher suites.
4. **Trust chain of certificate:** A valid chain of trust must be published by trusted certificate authorities. For this test, the certificate from the mail server is used to check if a valid chain of trust can be built.
5. **Public key of certificate:** Checks if the bit-length of the public key of the mail server certificate is sufficiently secure.
6. **Signature of certificate:** Checks if the signed fingerprint of the mail server certificate was created with a secure hashing algorithm.
7. **Domain name on certificate:** Checks if the domain name of the receiving mail server (MX) matches the domain name on the certificate.
8. **TLS compression:** Checks if the mail server supports TLS compression. The use of compression can give an attacker insights into secret parts of encrypted communication.

[2]Guidelines for TLS: https://www.ncsc.nl/actueel/whitepapers/ict-beveiligingsrichtlijnen-voor-transport-layer-security-tls.html

2.2.6 DNS-based Authentication of Named Entities (DANE)

TLS encryption is currently based on certificates issued by Certificate Authorities. Nowadays, there are hundreds of organizations acting as Certificate Authorities [22]. However, over the last few years, some serious security incidents have occurred regarding the certificate authorities. For example, some Certificate Authorities have been compromised and issued rogue certificates [10] [13].

DANE introduces a mechanism for domains to specify to clients which certificates should be trusted for the domain. DANE creates an explicit link where certificates are tied to a given domain. The domain owner creates a TLSA resource record in the DNS server, which identifies the certificate, the certificate authority or the key. When a client receives a certificate during the TLS negotiation, it looks up the TLSA record for that domain and matches the TLSA data against the certificate. This enables the client to verify if the certificate presented by the server is trusted/issued by the domain owner. DNSSEC is an prerequisite for DANE because DNSSEC creates a chain-of-trust and integrity of the TLSA data. An overview of DANE is displayed below [15] [30].

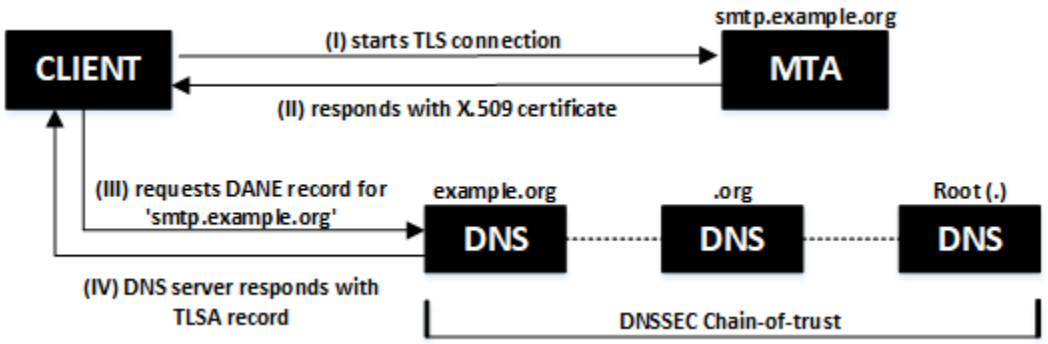

Figure 6: An overview of DANE.

The initial handshake takes place in plain text. This enables an attacker to conduct man-in-the-middle attack to make it appear that TLS is unavailable. This attack is called 'STRIPTLS attack'. DANE allows to advertise support for secure SMTP via a TLSA record. This tells connecting clients they should require TLS, thus preventing STRIPTLS attacks.

A DANE record contains a few parameters. The first parameter is the usage field. The other parameter is the selector. The selector has to be either 0 (Match full certificate) or 1 (Match only SubjectPublicKeyInfo). The last parameter is the matching type, which has to be 0 (Exact match on the selected content), 1 (SHA-256 hash of selected content) or 2 (SHA-512 hash of selected content)[24]. An example of a DANE record is displayed below. In this example, the usage field is 0, the selector is also 0 and the matching type is 1.

```
_25._tcp.mx1.example.org IN TLSA ( 0 0 1 <certificate data>
```

Figure 7: An overview of DANE.

Since the server certificate is directly specified by the TLSA record and a chain-of-trust has been established via DNSSEC, the certificate may be self signed. Note, that by creating a self-signed certificate, the 'Trust chain of certificate'-check for STARTTLS will result in a fail during the experiment.

Parameters that are being verified:

DANE allows certificates to be bound to DNS names using DNSSEC. Therefore, DANE will fail if the DNSSEC test has also failed. During the experiment, two parameters regarding DANE will be checked:

1. **DANE existence:** To check that each of the mail server domains provides a TLSA record for DANE.
2. **DANE validity:** To check if the DANE fingerprints presented by the mail server domains are valid for the mail server certificates.

3 Experiment

For this research, a data-set of Dutch organizations was created that was used as an input for the tool. This chapter discusses what type of information the data-set must contain in order to answer the research questions. Furthermore, this chapter also discusses how the data-set was scraped and parsed from the internet. In chapter 1, a list of email security techniques was defined that has to be verified during the experiment. The tool from 'internet.nl' was used for the experiment to verify if these email security techniques have been implemented.

3.1 Required information

The first step is to ensure that every Dutch organization is present in the data-set. A feasible solution is to use the data from the Dutch chamber of commerce (*Kamer van Koophande: KVK*). In the Netherlands registration in the commercial register from the Chamber of Commerce is compulsory for every organization. Each organization then gets a unique number called KVK-number. This KVK-number consists of 8 digits. The list of possible kvk-numbers (10^8 possibilities) will be used as a starting point for collecting a data-set of Dutch organizations [21].

To answer the sub-questions, each organization in the data-set must also contain additional information. To answer the first sub-question, a list of sectors is required. The Dutch Central Bureau for Statistics has defined a list to categorize the main economic activity of an organization. This list is called 'De Standaard Bedrijfsindeling' (SBI) and contains 21 categories [2]. Each category is split into multiple sub-categories. Both the main category and sub-category are included in the data-set[3].

To answer the second sub-question, the number of employees are required. The number of employees will be collected to determine the size of an organization. For this research, the number of employees will be split into five different groups. By splitting the number of employees into different groups, it would be easier to answer sub-question 2. Each of the 5 groups contains (1) 1-10 employees, (2) 11-50 employees, (3) 51-100 employees, (4) 101-250 employees or (5) '251 or more' employees.

To answer the third sub-question, a location is required. This means that every organization in the data-set should also contain the street name and zip code. A python library called Geopy[4] was used to retrieve additional information based on street name and zip code such as the coordinates of the location (which can be plotted on a map), the province and the municipality (*Dutch: Gemeente*) [12].

To summarize, a data-set of Dutch organizations should contain the KVK-number, Trading name, location, sector, employees and the domain name per organization in order to answer the research questions. The next paragraph discusses how the data for the data-set has been collected.

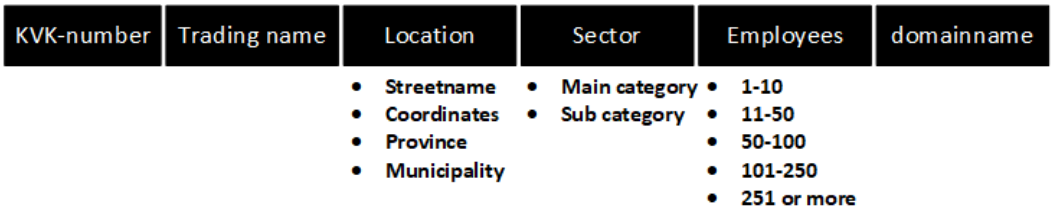

Figure 8: An overview of the data-set.

[3] A description for each sector can be found here: https://www.cbs.nl/nl-nl/onze-diensten/methoden/classificaties/activiteiten/sbi-2008-standaard-bedrijfsindeling-2008/de-structuur-van-de-sbi-2008-versie-2018

[4] https://geopy.readthedocs.io/en/stable/

3.2 Scraping and parsing the data-set

The previous paragraph described that every organization has to be registered at the chamber of commerce and has a unique KVK-number. This KVK-number consists of 8 digits. In theory, the total number of unique KVK-numbers contains 10^8 possibilities. The website of the chamber of commerce contains an API interface which can be queried to check if a KVK-number would exist. However, this is not a feasible solution due to the high number of possibilities. Fortunately, a list of existing KVK-numbers has been made available in a repository located on the internet.[5] This list contains 3.375.884 unique KVK-numbers.

A public website (SITE-1) [6] can be queried by submitting the KVK-number in the search field of the URL. The response from the website includes all the necessary data for the data-set (KVK-number, trade name, location, sector, employees and domain name). 3.375.884 unique KVK-numbers from the repository were reflected against the website (SITE-1). This resulted in 2.870.658 matches. During the scraping and parsing the data-set, it turned out that many organizations that have 1 up to 10 employees didn't contain a domain name. Most of these organizations didn't have a website or they only have a social media website (*such as Facebook*). Therefore, organizations that have 1 up to 10 employees are not added to the data-set. Of the 2.870.658 organizations that have been scraped and parsed from the website, only 58.749 organizations had more than 10 employees. Note, that 3.483 of the 58.749 organizations didn't contain a domain name.

Next, a second website (SITE-2)[7] was used to query the KVK-numbers that didn't have a record on the first website (SITE-1) by submitting the KVK-number along with the Trade name in a search engine[8]. The result from the search engine was the web page of the organization. Of the 505.226 KVK-numbers, only 6.960 organizations were present in the second website (SITE-2) that had more than 10 employees. This means that in total 65.709 unique organizations[9] have been parsed from both websites.

Unfortunately, the second website didn't contain a domain name of the organization. Therefore, a scaper was build to scrape the domain name from different search engines, where the streetname, zipcode and trading name were used in a search query the find the related domain name. The scraper couldn't retrieve the domain names for 4985 organizations. This means that only 60.724 organizations have been collected that contained a domain name. After analyzing the domain names, it turned out that different organizations share the same domain name. For example, franchises have a unique KVK-number, but might share the same domain name. Also, subsidiaries might share the same domain name. After filtering the duplicate domain names, a list of in total 50.521 organizations with unique domain names (along with KVK-number, trade name, location, sector and employees) has been scraped and parsed from the internet. This final data-set will be used as an input for the experiment.

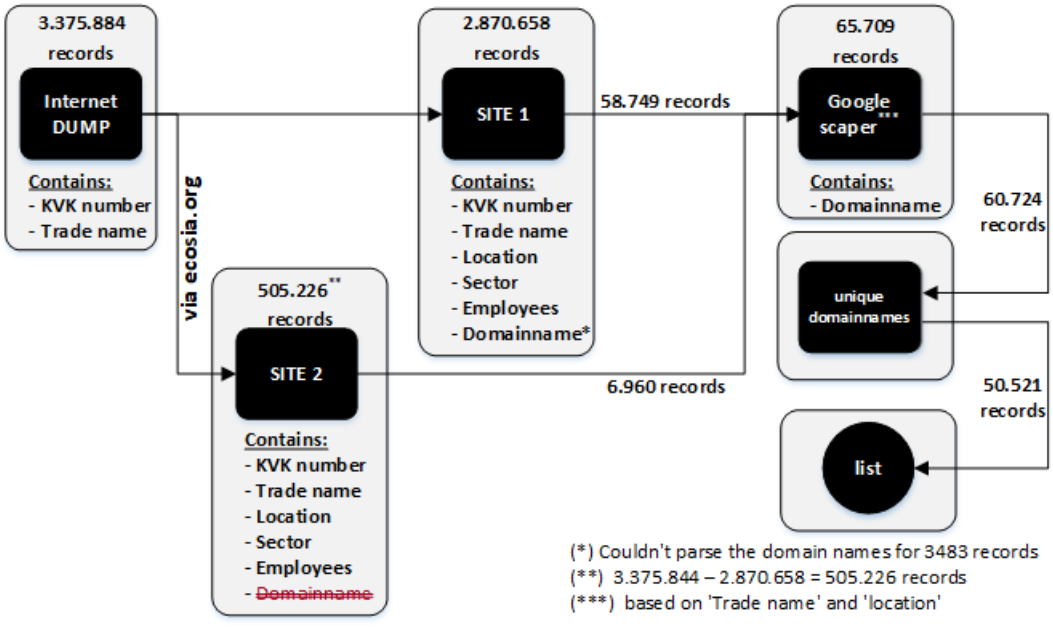

Figure 9: An overview of scarping and parsing workflow.

[5] https://github.com/digitalheir/kvk-handelsregister
[6] SITE-1: www.opencompanies.com
[7] SITE-2: www.ozoo.nl
[8] ecosia.org
[9] (SITE 1) 58.749 + (SITE 2) 6.960 = 65.709

3.3 Validation of the data-set

Before the data-set was used as the input for the experiment, some form of verification had to be done on the data-set to ensure that the collected data-set contains all the Dutch organizations. The collected data-set was reflected against the data-set that is available from the Dutch Central Bureau of Statics (CBS). The CBS keeps on behalf of the Dutch government record of every Dutch organizations. The CBS also separates the organizations based on the number of employees [10]. Note, that the data-set from the CBS only contains the total numbers of organizations and not the actual list of organizations.

Table 2 shows the accuracy of the collected data-set compared to the data-set from the CBS. The first row contains the number of organizations that exist in the Netherlands according to the CBS. The second row contains the number of organizations that have been collected through scraping and parsing. The third row contains the number of organizations that have been collected through scraping and parsing that contains a URL. The last row contains the number of organizations that have been scraped that contains unique URLs.

	11-50	51-100	101-250	251 +	Total
The number of organizations according to the CBS	51380	6895	4585	3120	65980
The number of organizations that have been collected through scraping	50668	7342	4519	3184	65709
The number of organizations that have been collected and contains a URL	46374	6991	4354	3005	60724
The number of organizations that have been collected with an unique URL	38748	5618	3612	2543	50521

Table 2: The total number of organizations divided by the groups with number of employees compared to the data-set from the CBS.

According to the data-set from the CBS, there are 65.980 organization with more than 10 employees.[11] In total, 65.709 organizations have been scraped. This is 99,59% accurate. However, there is a small deviation between the groups of employees. The groups with '51-100' and '251 and more' employees contain more organization compared to the data-set from the CBS, while the other groups contain fewer organizations. On the right side of figure 10 is the accuracy displayed per group in percentages. If you compare the number of organization per group with the number of employees, the smallest deviation is group '51-100' employees with 98.61%, and the largest deviation is group '11-50' employees with 90.26%. An overview of the number of organizations per group is displayed below[12].

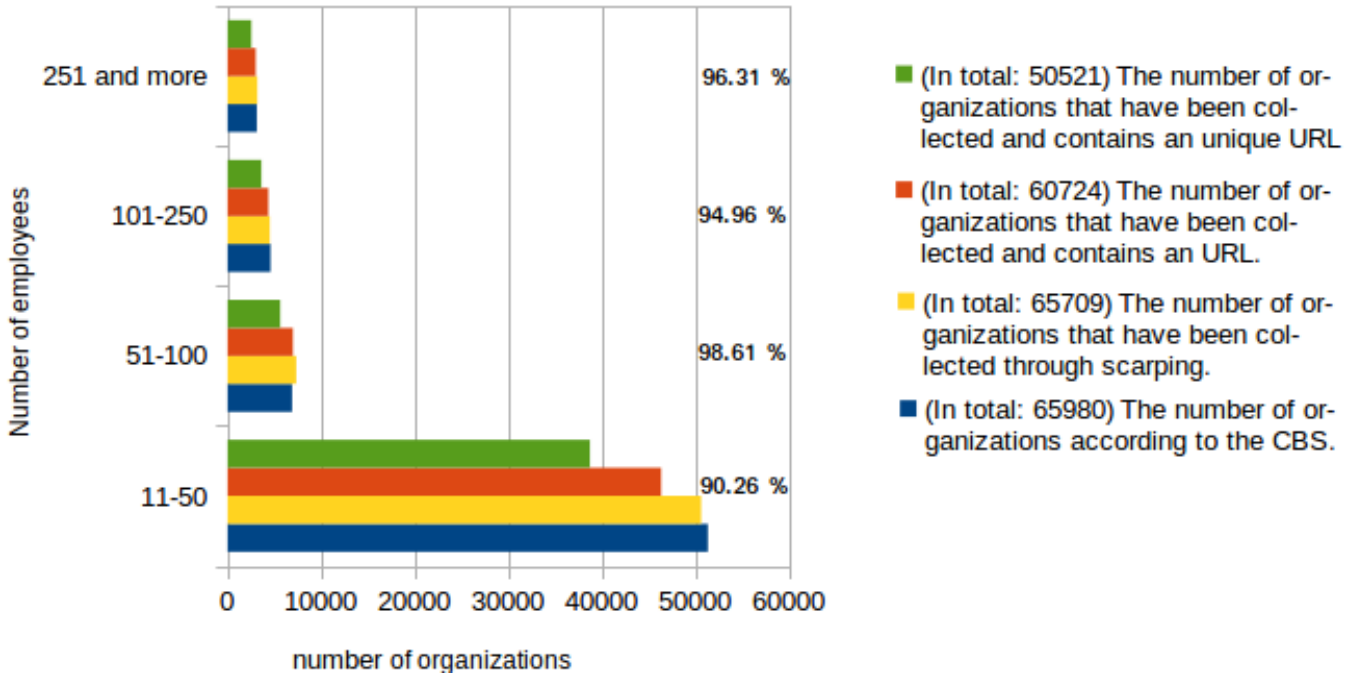

Figure 10: An overview of the collected data.

[10] https://opendata.cbs.nl/statline/CBS/nl/dataset/81589NED/table?ts=1529500476726
[11] The first quarter of 2018
[12] The collected data-set can be found in the appendix of this report.

3.4 Experimental setup

The tool from 'internet.nl' can verify if email security techniques have been adopted and only requires a domain name as input. The domain names from the collected data-set are used as the input. The tool queries the DNS server along with the SMTP server (listed as MX record in the DNS server). The tool contains a batch server with an API interface in which all the 50.521 domain names can be submitted at the same time. The actual experiment took approximately 4 days to complete and the output of the experiment was available in json format. Another parser was built to extract the information from the json file and the information was added to the data-set for further analysis.

3.5 Remarks about the experiment

There are a few remarks about the collected data-set and the tool that was used:

Remarks about the data-set

There are a few remarks about the collected data-set:

1. The data-set from the CBS only contains aggregated numbers. Therefore, it is not possible to check if an organization from the collected data-set is also present in the data-set from the CBS. For example, the data-set from CBS might contain an organization that is not present in the collected data-set.
2. 4.985 organizations did not contain a domain name. It might happen some of these domain names actually exists.
3. Based on figure 10, a relatively large number of organizations that contain 11-50 employees have the same domain name. After analyzing these organizations it turned out that many of these organizations are franchises that share the same domain name (*such as Albert Heijn, Jumbo supermarkten and Rabobank)*.
4. The repository that was used did not contain very accurate information since the repository was created in 2015. This means that some *(new)* organizations are not present in the repository.
5. The created data-set didn't contain organizations with the number of employees between 1 and 10. Most of these organizations didn't have a website or they only have a social media website (*such as Facebook*). Therefore, organizations that have 1 up to 10 employees are not added to the data-set.

Remarks about the tool

There are a few remarks regarding the experiment. The remarks are listed below.

1. The tool could not retrieve the MX records for 3871 domains. Therefore, 46.650 out of 50.521 domain names have been used in the experiment. Mail servers might use the A or AAAA record in case an MX record does not exist. However, this fallback is considered as legacy. Therefore, the tool specifically requires an MX record.

2. The tool has a limitation. The tool could only check if a DNS server contains a domainkey sub-record. In order to completely check whether DKIM has properly been implemented, you need to receive an email from the sending MTU that signs the mail with the private DKIM key to completely verify DKIM. There is no feasible solution to query and evaluate the public key in your DKIM record, because the DKIM selector (that should be in the emails you send) is needed to do so.

4 Results

This chapter discusses the results of the experiments and answers the research questions. The 4 research questions have been addressed in subsection 1.1. Each research question will be answered and discussed in the next paragraphs.

Question 1: How many email security techniques have been adopted?

The tool from the experiment have checked 19 different parameters for 46.650 organizations. The results for each parameters is displayed below in percentages. The average adoption rate for the 19 different parameters is 45,58 %. This means that from the 19 different parameters, only 8,66 parameters have been adopted.

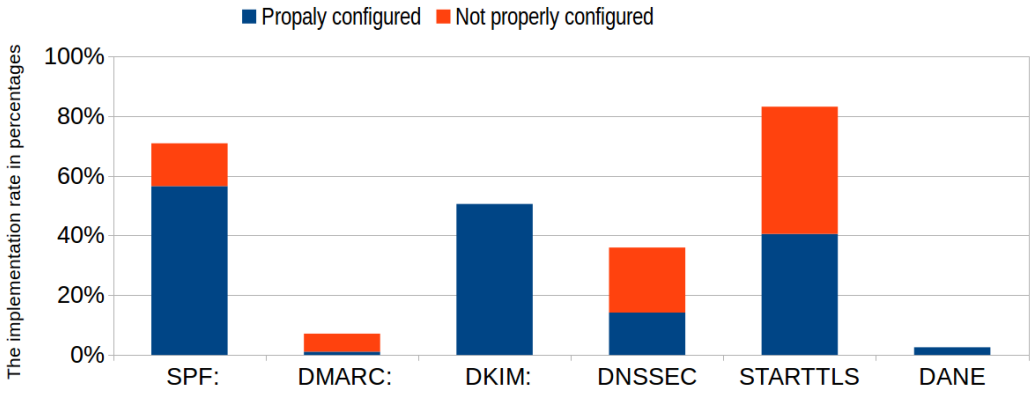

Figure 11: An overview of sufficient and insufficient adopted email security techniques.

Not every organization uses a sufficient policy. 79,77 % of the adopted SPF policy are considered sufficient (hard-fail or soft-fail), while only 15,47 % of the DMARC policies are considered sufficient (reject or quarantine policy). This means that not every organization have properly configured these techniques. Furthermore, for organizations that have adopted DNSSEC, only 39,45 % pass the checks for all the DNSSEC parameters (DNSSEC valid, DNSSEC MX exists and DNSSEC MX valid). For the organizations that have adopted STARTTLS, less than 50 % percent (48,73 %) uses all the parameters and settings according to the guidelines from the NCSC [23]. Only a single organization that has a DANE record didn't pass the DANE validation check. Note, the experiment wasn't able to determine if DKIM is valid (see: subsection: 2.2.2). Figure 12 graph shows that newer techniques such as DANE and DMARC have a lower adoption rate compared to older techniques such as SPF.

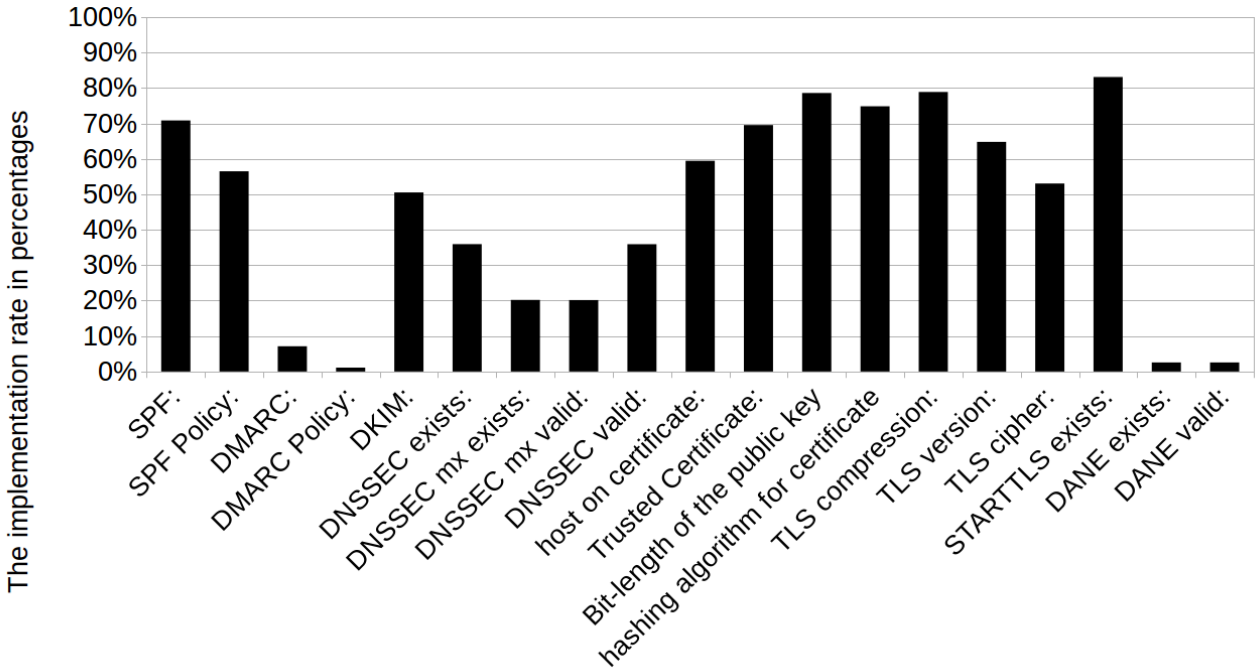

Figure 12: An overview of adopted email security techniques for Dutch organization: The adoption rate is displayed in percentages per technique for the 46.650 organizations.

Question 2: Is there a distinction between the size of organizations regarding the adoption of email security techniques?

The number of employees have been split into four different groups. The adoption rate of email security techniques has been determined for each group by dividing the total number of organizations by the total number of adopted techniques, which makes it possible to determine whether or not there was a distinction between small, medium and large organizations regarding the adoption of email security techniques.

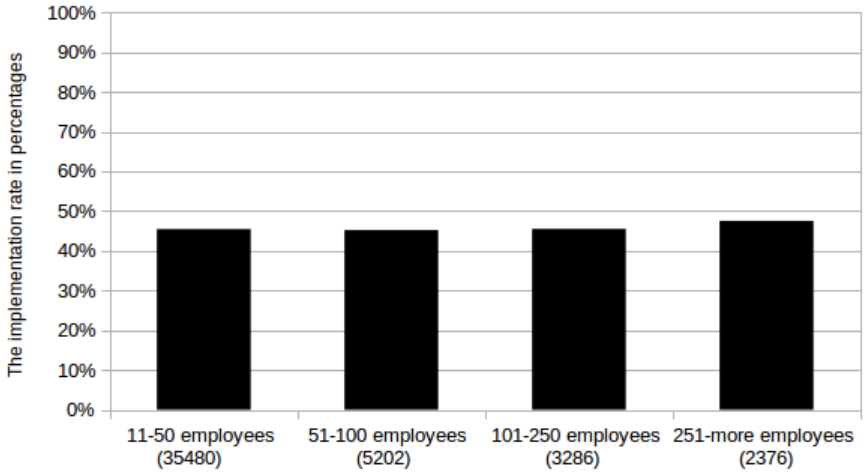

Figure 13: The adoption rate of email security techniques divided by the number of employees.

Figure 14 displays the adoption rate per technique for each group in percentages. SPF is slightly less adopted by smaller organizations. Only the group with more than 251 employees has an adoption rate that is 2,3 % higher compared to the other groups. In conclusion, there no distinction between the number of employees regarding the adoption of email security techniques.

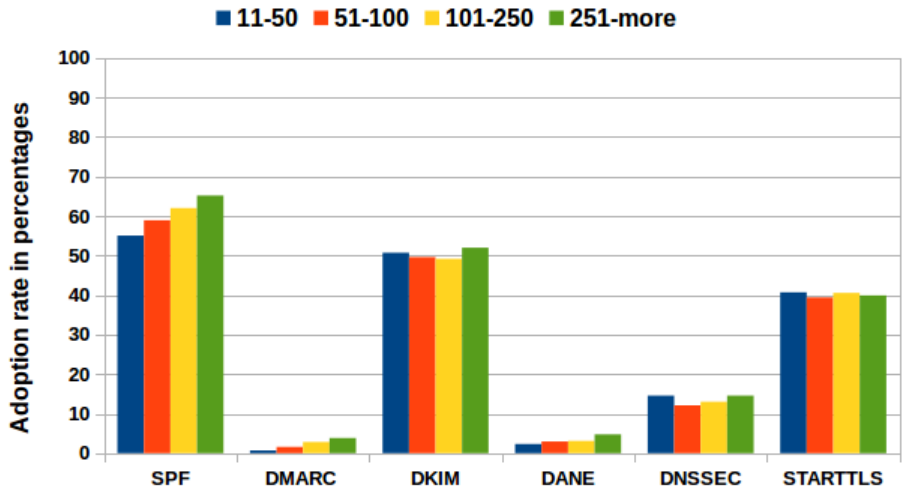

Figure 14: The adoption rate per email security techniques divided by the number of employees.

Question 3: Is there a geographical distinction between organizations regarding the adoption of email security techniques?

Two graphs have been created to answer this question. The first graph shows an overview of the adoption rate per province. The second graph shows an overview of the adoption rate per municipality (*Dutch: 'Gemeenten'*). Based on figure 15, it can be concluded that there is almost no distinction between the adoption rate between provinces. The difference between province with the highest and lowest adoption rate is only 3,0%.

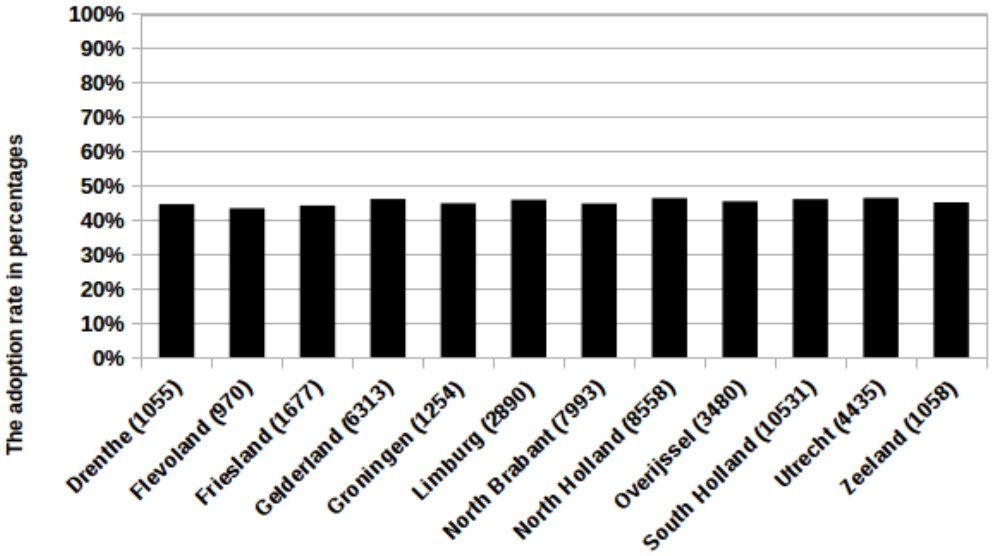

Figure 15: The adoption rate of email security techniques divided per province.

A heatmap per municipality was also created to determine if there is a geographical distinction on a more detailed level compared to figure 15, which was based on provinces. 350 municipalities have been plotted on the heatmap in figure 16. The average score of adopted techniques per municipality is displayed on the heatmap. The average score has been calculated by the sum of the total score divided by the number of organizations. Note, that the heatmap only displays the average adoption score per municipality. It might happen that there are only a few organizations present in a municipality and therefore strongly influence the average score.

The minimum average score is 6 and the maximum is 10. Only 1 municipality has an average score of 6, 12 municipalities have an average score of 7, 117 municipalities have an average score of 8, 195 municipalities have an average score of 9 and 25 municipalities have an average score of 10. The municipality that has an average score of 6 is Gemeente Echt-Susteren from the province Limburg. This municipality contains 4 organizations of which 3 organizations have more than 50 employees. The other organization contains 11 employees.

Based on figure 15 and figure 16, we can conclude that there is no geographical distinction regarding the adopted email security techniques.

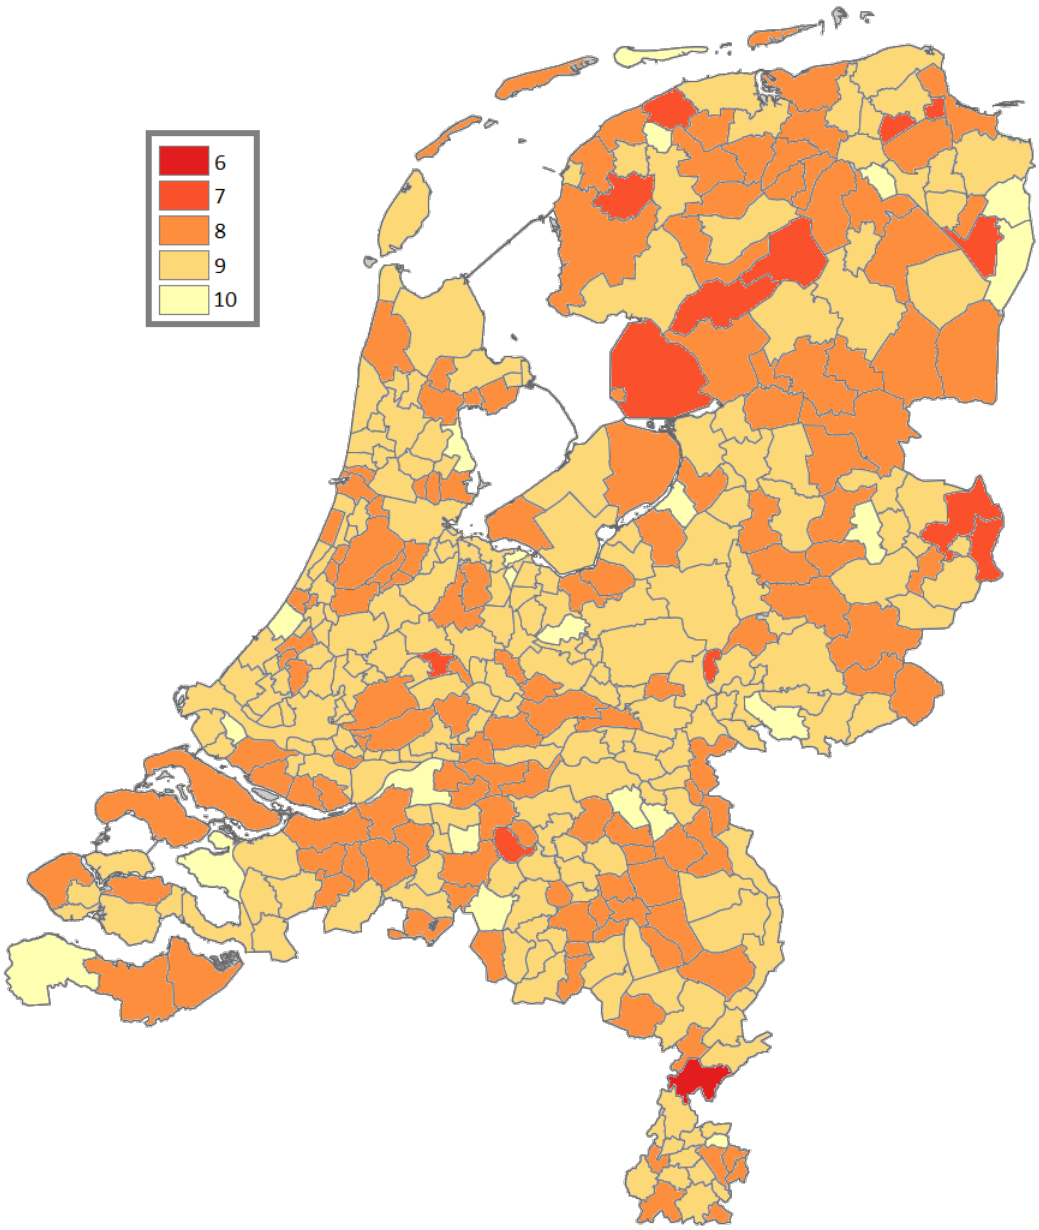

Figure 16: A heatmap of the adoption rate of email security techniques per municipality.

Question 4: Is there a distinction between the different sectors regarding the adoption of email security techniques?

Every organization is present in one of the 21 different types of sectors. The average adoption score for each sector has been plotted on the graph below [13].

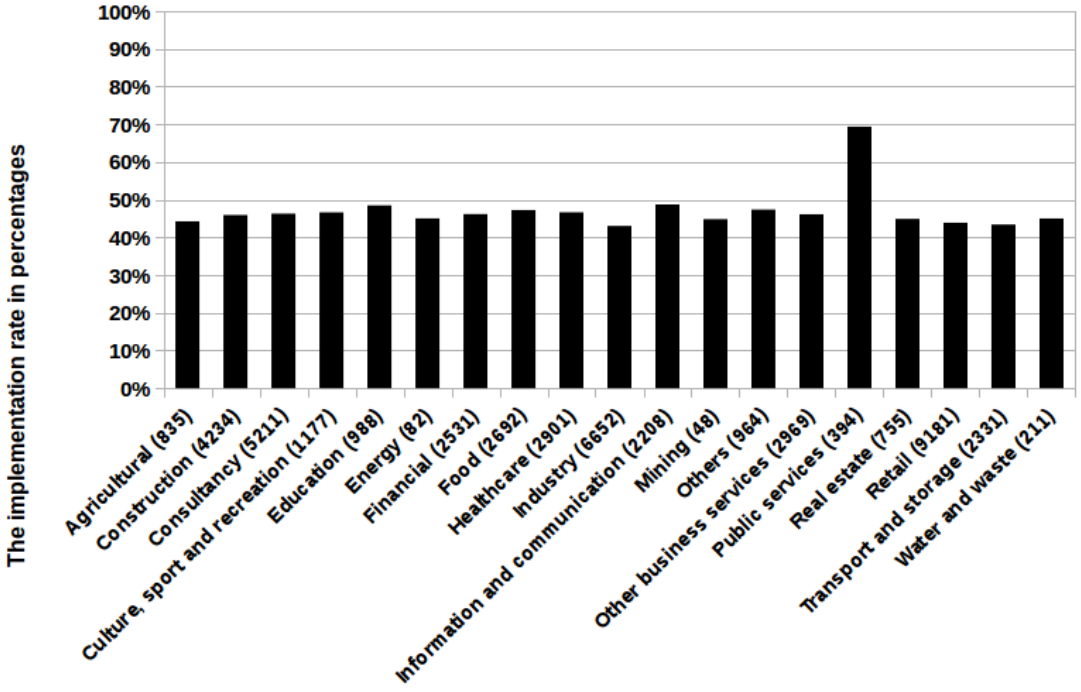

Figure 17: The adoption rate of email security techniques divided per sector. Each sector also contains the number of organizations that are active in the sector.

17 sectors score between 8,41 and 9,21. The sector with the highest score is the 'Public services' sector with an average score of 13.18. It is most likely that the difference might be due to the fact that the Dutch government has made compulsory policies for governmental organizations to implement SPF, DKIM, DMARC, STARTTLS and DNSSEC. Many governmental organizations are present in the 'Public services' sector. No other explanation has been found why the 'Public services' sector has scored higher compares to other sectors. Therefore, the assumption is that the high score is related to compulsory policies from the Dutch government. [24].

Furthermore, the adoption rate of each email security techniques has also been investigated. A graph for the SPF policy technique and the DKIM technique has been created. The SPF and DKIM graphs show no interesting remarks besides the fact the 'Public Services' sector has the highest score.

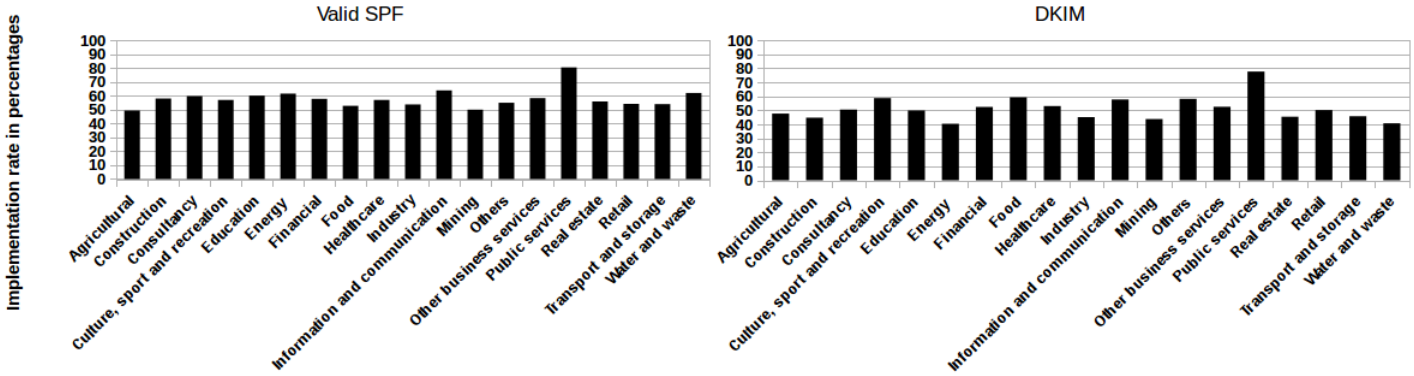

Figure 18: An overview of the adoption rate for SPF policy and DKIM per sector displayed in percentages.

[13] A description for each sector can be found here: https://www.cbs.nl/nl-nl/onze-diensten/methoden/classificaties/activiteiten/sbi-2008-standaard-bedrijfsindeling-2008/de-structuur-van-de-sbi-2008-versie-2018

A graph for DMARC and DANE was also created. The DMARC graph clearly shows that organizations in the 'Public Services' sector have the highest score. However, the adoption rate for the 'Public Services' sector is still less than 20 %. The DANE graph shows that no organization in the mining sector has properly adopted DANE.

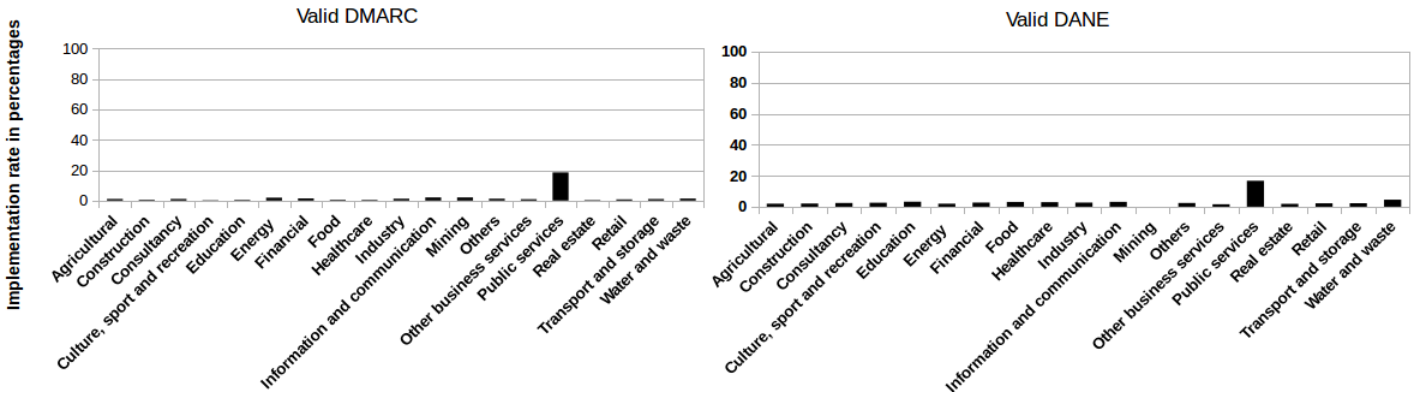

Figure 19: An overview of the adoption rate for DMARC policy and DANE per sector displayed in percentages

The graphs for DNSSEC and STARTTLS are displayed below. Note, that only organizations that have properly adopted DNSSEC and STARTTLS are displayed in the graphs. For example, if DNSSEC didn't contain a signed MX record, the organization is not considered to have properly adopted DNSSEC and therefore hasn't been added to the graph. Just like the previous graphs, the 'Public Services' sector has the highest score for DNSSEC. However, two other sectors have a higher score regarding the STARTTLS graph, which are the 'Education' sector and the 'Information and Communication' sector. No explanation has been found why the 'Education' sector and the 'Information and Communication' sector have a high score.

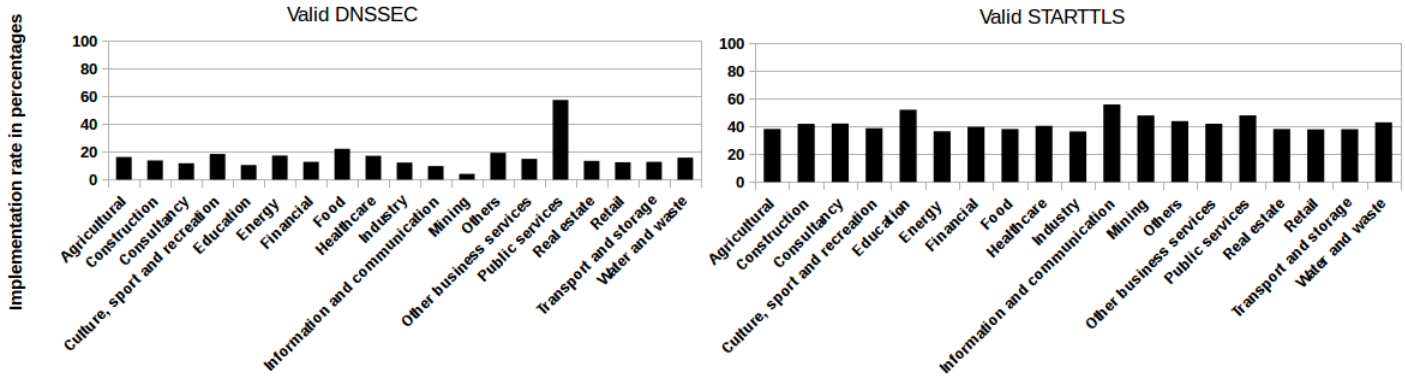

Figure 20: An overview of the adoption rate for DNSSEC and STARTTLS per sector displayed in percentages

4.1 Interesting findings

Some interesting findings have been found during the analysis of the results. This paragraph briefly discusses these findings:

1. The top 1000 organizations with the most employees score an average of 9,30. That is 3,37 % higher compared to the average score of all the organizations.
2. The Amsterdam Exchange Index (AEX) index is one of the most important indicator of the Dutch stock market. Organizations that are present in the AEX index score an average of 10,32. That is 8.34% higher compared to the average score of all the organizations.
3. The list of results only contains 4 organizations that didn't pass the 'TLS trusted certificate' check but did contain a valid DANE record along with a valid DNSSEC MX record.
4. The subsector that has the lowest score is the 'Manufacture of aircraft parts' subsector with an average score of 3,2. Interestingly, some of the organizations have contracts with the Dutch Military.

5 Discussion and Conclusion

This chapter discusses the results, conclusion and the future work.

5.1 Discussion

We have already briefly discussed some remarks about the research in previous chapters. This section will discuss the complete project along with the remarks about the results.

Remarks about the results:
There are a few remarks regarding the results. The list of the remarks is displayed below.

1. 8 of the 19 parameters that were checked are related to STARTTLS. This means that when an organization has adopted every technique except STARTTLS, it only get 11 points.
2. There might be only a few organizations present in a municipality and therefore strongly influence the average score for the municipality.
3. The 'Public services' sector has scored higher compared to other sectors. The assumption is that the high score is related to compulsory policies from the Dutch government because many governmental organizations are present in the 'Public services' sector.

5.2 Conclusion

This research investigated which and how many email security techniques have been adopted by organizations within the Netherlands. A list of Dutch organizations with more than 10 employees (in total 46.650 unique organizations) has been created and we defined 19 different parameters that were checked during the experiment. Next, we used the tool from 'internet.nl' to check whether or not these email security techniques have been adopted.

Based on the results it turned out that an organization has on average adopted 8,66 of the total of 19 parameters. This means that email servers from Dutch organizations have adopted less than 50 % of the email security techniques that have been defined by the Dutch government. We didn't find a relation between the the size of an organization or the geographical location in term of the adoption rate. However, we did find a relation between the type of sector. The 'Public Services' sector has the highest score with an average score of 13,18. We assume that the high score is related to compulsory policies for governmental organizations because many governmental organizations are present in the 'Public services' sector.

5.3 Future Work

First of all, we propose to investigate whether or not there is a distinction between hosting providers. Some hosting providers might have adopted fewer email security techniques compared to other hosting providers. Therefore, additional research have to be done. Secondly, we propose to create data-set of Dutch organizations that is more up-to-date. Thirdly, new organizations will also be measured. Furthermore, we propose to repeat the experiments in the future to determine if the adoption rate has been increased. For example, we advise repeating the experiment periodically.

References

[1] Bart knubben, forum standaardisatie, lijst open standaarden, jun 2012. https://www.forumstandaardisatie.nl/open-standaarden/lijst/verplicht?f

[2] Centraal bureau voor de statistiek, standaard bedrijfsindeling 2008, version 2018, jan 2018. https://www.cbs.nl/-/media/$_{pdf/2018/17/sbi\quad 2008\quad versie\quad 2018.pdf}$.

[3] Centraal bureau voor de statistiek, statsline, bedrijven bedrijfsgrootte en rechtsvorm, 22 january 2018. http://statline.cbs.nl/StatWeb/publication/?DM=SLNLPA=81588ned.

[4] cloudflare, how dnssec works, cloudflare, may 2018. https://www.cloudflare.com/dns/dnssec/how-dnssec-works/.

[5] D. crocker, domainkeys identified mail (dkim), jan 2018. http://www.dkim.org/info/DKIM-teaser-03.pdf.

[6] Dmarc one year later, terry zink: Security talk, december 3, 2015. https://blogs.msdn.microsoft.com/tzink/2015/12/03/dmarc-one-year-later-and-what-have-we-learned.

[7] Dmarc.org, overview, jan 2018. https://dmarc.org/overview/.

[8] Dovecot, mailserveroverview, dec 2012. https://wiki.dovecot.org/MailServerOverview.

[9] E. zwicky, ed. yahoo, domain-based message authentication, request for comments: 7489, mar 2015. https://tools.ietf.org/html/rfc7489.

[10] European union agency for network and information security (enisa), certificate authorities - the weak link of internet security, sep 2016. https://www.enisa.europa.eu/publications/info-notes/certificate-authorities-the-weak-link-of-internet-security.

[11] fastmail, ssl vs tls vs starttls, 2018. https://www.fastmail.com/help/technical/ssltlsstarttls.htm.

[12] geopy, geopyâs documentation, version 2018, jan 2018. https://geopy.readthedocs.io/en/stable/.

[13] Govcert.nl, frauduleus uitgegeven beveiligingscertificaat ontdekt, factsheet fs 2011-06, sep 2011. https://geopy.readthedocs.io/en/stable/.

[14] Improving email security, richard c, 15 sep 2017. https://www.ncsc.gov.uk/blog-post/improving-email-security.

[15] Internet engineering task force (ietf), rfc 6376, domainkeys identified mail (dkim) signatures, sep 2011. https://tools.ietf.org/html/rfc6376.

[16] Internet engineering task force (ietf), rfc 7208 sender policy framework (spf) for authorizing use of domains in email, version 1 ,2014. https://tools.ietf.org/html/rfc7208.

[17] Internet-wide efforts to fight email phishing are working, elie bursztein and vijay eranti, 6 december 2013. security.googleblog.com/2013/12/internet-wide-efforts-to-fight-email.html.

[18] Internet-wide efforts to fight email phishing are working, google security blog, december 6, 2013. https://security.googleblog.com/2013/12/internet-wide-efforts-to-fight-email.html.

[19] Internet.nl now also checks strictness anti-mail-spoofing standards, internet.nl, january 9, 2018. https://internet.nl/about/.

[20] Julian mehnle, openspf,sender policy framework, modifier/redirect, feb 2009. http://www.openspf.org/Modifier/redirect.

[21] Kamer van koophandel, nummers in het handelsregister, may 2018. https://www.kvk.nl/over-de-kvk/over-het-handelsregister/wat-staat-er-in-het-handelsregister/nummers-in-het-handelsregister/.

[22] Mozilla included ca certificate list, certificate authorities mozilla, june 2018. https://ccadb-public.secure.force.com/mozilla/IncludedCACertificateReport.

[23] Nationaal cyber security centrum, ict-beveiligingsrichtlijnen voor transport layer security (tls), nov 2016. https://www.ncsc.nl/actueel/whitepapers/ict-beveiligingsrichtlijnen-voor-transport-layer-security-tls.html.

[24] National institute of standards and technology, trustworthy email,nist special publication 800-177, sep 2016. https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-177.pdf.

[25] Network working group, dns security introduction and requirements,request for comments: 4033, 2005. https://tools.ietf.org/html/rfc4033.

[26] Network working group, smtp service extension for secure smtp over transport layer security, request for comments: 3207, feb 2002. https://tools.ietf.org/html/rfc3207.

[27] Nieuwe adoptieafspraak voor informatieveiligheidstandaarden , overheidsbrede beleidsoverleg digitale overheid, 18 april 2018. https://www.forumstandaardisatie.nl/open-standaarden/lijst/verplicht.

[28] Philip hazel, exim: The mail transfer agent, o'reilly media, inc., feb 2001. ISBN: 9780596000981.

[29] postmarkapp.com, what is dmarc?, request for comments: 7489, may 2018. https://postmarkapp.com/support/article/892-what-is-dmarc.

[30] Richard blum; christine bresnahan, lpic-2: Linux professional institute certification study guide, 2nd edition,john wiley & sons, 2016. ISBN: 9781119150794.

[31] Technical note (nist tn) - 1945 , nist: National institute of standards and technology, feb 2017. https://www.nist.gov/publications/email-authentication-mechanisms-dmarc-spf-and-dkimt.

[32] Trust anchor, dnssec-tools wiki systems, april 2008. https://www.dnssec-tools.org/wiki/index.php?title=Trust$_{Anchor}$.

6 Appendix

1. **DAT-SET:** The data-set has not been published due to the fact that someone (such as spammers) can abuse the data in the data-set. The collected data-set can be retrieved by contacting the author (vincent.vandongen**@**os3.nl)