jamk

Enforcing SMTP encryption

Research on current DANE and MTA-STS implementations in Finland

Jukka Nopanen

Master's thesis
May 2022
School of Technology
Master's Degree Programme in Information Technology, Cyber Security

jamk | Jyväskylän ammattikorkeakoulu
University of Applied Sciences

jamk

Description

Nopanen, Jukka

Enforcing SMTP encryption
- Research on current DANE and MTA-STS implementations in Finland

Jyväskylä: JAMK University of Applied Sciences, May 2022, 64 pages.

Technology, Cyber Security. Degree Programme in Information Technology. Master's thesis.

Permission for open access publication: Yes

Language of publication: English

Abstract

SMTP-protocol, that is used to transmit email messages, was originally defined in 1982. Original protocol version was defined during the early years of Internet, when the threats against email communications were very much different from the current ones, and the security of communications was based on mutual trust and limited access to the network. The protocol was plain-text and did not support encryption. Development of technology, globalization of Internet and explosive spread of usage of email in all aspects of society has changed old threats and created new ones. New protocols have been defined to protect from these threats, that improves confidentiality of email communications without actions from the end users.

The purpose of the study was to research how widely these protocols were used at the time of making the study in the Finnish country code top-level domain .fi. In addition to that, a documentary analysis was conducted to find international publications and statistics about the usage and implementation rate of the protocols.

The study was conducted by researching the implementation rate of DNSSEC, DANE, MTA-STS and TLS-RPT and the validity of the settings in the .fi ccTLD namespace. Information was gathered by executing multiple rounds of DNS and HTTP queries about addresses specified in the standards. Responses were saved, analyzed, and classified according to their validity. In addition, a document search was made to find international recommendations and statistics.

Results of the literature search show that some international recommendations exist, but a comprehensive recommendation is widely missing. Comprehensive information about the implementation rate was not found. According to the results from .fi domain, implementation rate of protocols discussed in the study was very low, even close to non-existent.

Results of the research show that when implemented correctly, these protocols improved confidentiality, but there was a lot to improve on the implementation rates.

Keywords/tags (subjects)

Finland, DNSSEC, DANE, MTA-STS, TLS-RPT, implementation rate

Miscellaneous (Confidential information)

jamk

Kuvailulehti

Nopanen, Jukka

SMTP-liikenteen salauksen pakotus
- Tutkimus tämänhetkisistä DANE ja MTA-STS toteutuksista Suomessa

Jyväskylä: Jyväskylän ammattikorkeakoulu. Toukokuu 2022, 64 sivua.

Tekniikan ala. Kyberturvallisuuden tutkinto-ohjelma. Opinnäytetyö, YAMK.

Julkaisulupa avoimessa verkossa: kyllä

Julkaisun kieli: englanti

Tiivistelmä

Sähköpostin välitykseen käytettävä SMTP-protokolla on määritelty alun perin 1982. Alkuperäinen protokollaversio on määritetty Internetin alkuaikoina, jolloin uhkakuvat poikkesivat merkittävästi nykyisistä, ja viestien turvallisuus perustui yleiseen luottamukseen ja verkon rajattuun pääsyyn. Protokolla oli selväkielinen, eikä tukenut salausta. Teknologian kehittyminen, Internetin kansainvälistyminen, ja sähköpostin käytön räjähdysmäinen laajeneminen kaikille yhteiskunnan osa-alueille ovat muuttaneet vanhoja ja luoneet uusia uhkakuvia. Näitä torjumaan on kehitetty uusia protokollia, joita käyttämällä voidaan parantaa sähköpostin luottamuksellisuutta ilman loppukäyttäjien toimenpiteitä.

Tutkimuksen tarkoituksena oli selvittää, kuinka laajasti näitä protokollia käytettiin tutkimuksen tekohetkellä Suomen maakohtaisessa tunnisteessa, eli .fi-osoiteavaruudessa. Lisäksi tutkittiin kansainvälisiä ohjeistuksia ja protokollien käyttöastetta ja laadittiin pohja kansalliselle suositukselle protokollien käytöstä.

Tutkimus toteutettiin kartoittamalla DNSSEC:n, DANE:n, MTA-STS:n ja TLS-RPT:n käyttöastetta ja asetusten laatua koko .fi-osoiteavaruudessa suorittaen useita nimipalvelutietuekyselykierroksia ja HTTPS-pyyntöjä standardin mukaisiin osoitteisiin. Vastaukset tallennettiin, analysoitiin, ja luokiteltiin niiden oikeellisuuden perusteella. Lisäksi tehtiin hakuja kansainvälisten suositusten ja tilastojen löytämiseksi.

Tulokset osoittivat, että aiheeseen liittyviä kansainvälisiä ohjeita löytyy jonkin verran, mutta kattava ohjeistus puuttuu monilta osin. Kattavaa tietoa kansainvälisestä käyttöasteesta ei löytynyt. Suomen toteutustason osalta tulokset osoittavat, että tutkittujen tekniikoiden käyttöaste on erittäin matala, joiltakin osin jopa olematon.

Tutkimuksen tulokset osoittivat, että oikein toteutettuna protokollat parantavat sähköpostin luottamuksellisuuden turvaa, mutta niiden käyttöasteen kasvattamisessa oli merkittävästi parannettavaa.

Avainsanat (asiasanat)

Suomi, DNSSEC, DANE, MTA-STS, TLS-RPT, käyttöaste

Muut tiedot (salassa pidettävät liitteet)

Contents

Figures

Tables

Acronyms

API	Application Programming Interface
CA	Certificate Authority
ccTLD	Country code top-level domain
DANE	DNS-based Authentication of Named Entities
DKIM	DomainKeys Identified Mail
DMARC	Domain-based Message Authentication, Reporting and Conformance
DNS	Domain Name System
DNSSEC	Domain Name System Security Extensions
DS	Delegation Signer
HTTP	Hypertext Transfer Protocol
HTTPS	Hyper Text Transfer Protocol Secure
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IMAP	Internet Message Access Protocol
MSA	Mail Submission Agent
MTA	Mail Transfer Agent
MTA-STS	MTA Strict Transport Security
MUA	Mail User Agent
MX	Mail Exchange
PGP	Pretty Good Privacy
POP	Post Office Protocol
PTR	Pointer
RFC	Request for Comments

RR	Resource Record
S/MIME	Secure/Multipurpose Internet Mail Extensions
SMTP	Simple Mail Transfer Protocol
SPF	Sender Policy Framework
TCP	Transmission Control Protocol
TLS	Transport Layer Security
UDP	User Datagram Protocol

1 Introduction

Email has become one of the most, if not the most, important communications method between companies, organizations, governments, and individuals since it's invention. In the early days it was mostly used by the academic research society and confidentiality was not the top priority. Security was based on limited accessibility and mutual trust. Since that our society has had great advances in digitalization, and nowadays much of the everyday communications, including official communications, is done via email. Statista Research Department estimates that there are 4.3 billion email users, and 333 billion email messages send and received daily in 2022 (Statista, 2022).

Many communications include confidential information, like classified business information, personally identifiable information, banking information and health information. Different threat actors, ranging from criminals to nation states have discovered the possibilities of eavesdropping different communications methods (JRC, 2015).

Main findings of the JRC 2015 report are:

- Email communications are generally not sufficiently protected
- There are standards, protocols, and techniques capable of enhancing the security of email communications, but they are not always used or implemented properly in practice
- Mature and interoperable end-to-end email security solutions exist but are rarely used in practice
- Email communication channels (SMTP to SMTP) are not sufficiently protected in practice
- Lack of security in DNS has a direct impact on the security of email communications
- Email identity spoofing is still a major risk in email communications

During the last years more focus has been put on how to protect email from forgery as different attacks utilizing forged sender information have emerged (for example phishing). Different technologies such as Sender Policy Framework (SPF), DomainKeys Identified Mail (DKIM) and Domain Message Authentication (DMARC) have been developed to protect the integrity of email's, especially the sender information.

Less attention has been put on securing the confidentiality of emails. Transferring emails rely on a protocol called Simple Mail Transfer Protocol (SMTP), that was originally drafted in 1982 (Postel, 1982). It defines how emails are sent between different mail servers and originally it doesn't include support for encryption. However nowadays there are several options available to use encryption at the application layer, but that usually requires some extra effort from the end user, like installing an ad-on to the mail client and exchanging encryption keys with the recipients. Such things are often seen as too difficult use in normal everyday communications.

One of the suggestions in the JRC report was to promote security of email communication channels. Two new techniques have been developed to increase the confidentiality of SMTP connections between mail servers: DNS-Based Authentication of Named Entities (DANE) (IETF, 2015a) and SMTP MTA Strict Transport Security (MTA-STS) (IETF, 2016). No research has been done so far on the adoption rate of these technologies in the Finnish country-code top-level domain *.fi*.

This Thesis work was ordered by the National Cyber Security Centre of Finland (NCSC-FI). It has the responsibility of supporting, guiding, and monitoring information security in electronic communications.

With the increasing support of these technologies by software and service providers and the results of this thesis work NCSC-FI is hopefully able to improve the adoption level of these protective measures by increasing awareness and providing technical guidelines to domain operators.

2 Research framework

2.1 Scope of research

There are previous studies about preventing email forgery in *.fi* country top-level domain and specifically in the public sector in Finland (Kontinen, 2020) (Turunen, 2021). Conclusion of these studies is that implementation level of technical measures to prevent email forgery is still low.

Neither of the studies take a deeper look on how to protect the confidentiality of email communications as SMTP is originally a plain text protocol and Kontinen proposes this a topic for future research.

Based on the proposal and background research this was selected to be the scope for this research.

2.2 Study goals

The goals for this study can be summarized in two research questions:

1. How to deploy DANE and MTA-STS securely?
2. What is the deployment rate of them in the *.fi* country code top-level domain (ccTLD)?

2.2.1 Best practices

To find an answer to this research question, the study aims on doing a document analysis about the related technologies, search for best practices and recommendations from different authorities and organizations and summarizing the results as what could be used as a basis for national recommendations and requirements in Finland by NCSC-FI.

2.2.2 Implementation status in fi-domains

The purpose of this research question is to find out how widely these technologies are adopted in the *.fi* ccTLD and to help NCSC-FI to direct resources into correct target groups to increase awareness and improve the adoption rate.

2.3 Research methods

The research questions in this study are relatively simple. Security controls researched are already standardized, but relatively new. Standards and guidelines are available online, but no information

about the implementation rate in Finland isn't available. Based on this, qualitative, and especially document analysis, was chosen as the research method for this study. The study was divided into four phases:

1. Planning
2. Document analysis
3. Information gathering
4. Analysis of gathered data

In the first phase the topic and scope of the study was agreed with NCSC-FI and research permission to get required material was given.

The second phase is done as a document analysis, where already existing standards, documents and statistics are used as the source of information. The purpose of this part is to find and investigate information and research data already available about the subject to understand and evaluate the protocols researched, and to design tools used for information gathering and analysis. In this study it was done by searching the Internet for standards, recommendations, and statistics about the implementation of the covered protocols.

The third phase was done by measuring how many *.fi* domains have properly implemented the protocols covered in this study. The data gathered in this phase could be considered as a quantitative research, but the aim was not to build any numerical models from them but rather study the results in qualitative manner to see the implementation rate and if the implementations are securely done.

As the number of domains in the scope of this study is large, information collection and analysis had to be automated. Methods for information gathering are described in chapter 2.5.

In the last phase, gathered information was analyzed, and statistics were formed based on existence of valid deployment of the studied protocols. All corner cases were manually verified, and some additional random manual checks were done to verify the results.

2.4 Technical operation

Key technologies related to this study are DNS, SMTP, STARTTLS, DNSSEC, DANE, MTA-STS and TLS-RPT. They are all explained in more detail on the following sub-chapters.

2.4.1 Domain Name Service (DNS)

Routing of emails is based on mail exchanger (MX) resource records (RR) published in the Domain Name Service (DNS) as specified in RFC 1035 (Mockapetris P, 1987). Sending SMTP server performs a DNS MX query to the recipient address to find what SMTP servers are accepting mail on behalf of the recipient address. If a domain does not have MX records published, the sending SMTP server will try to send the message to the IP address of the domain itself, if available. If a domain has multiple MX records, the one with the lowest priority is attempted first and if it fails, one with a lower priority is attempted. Sample MX record can be seen below:

```
domain.fi.    3600    IN      MX      10 smtp1.domain.fi.
domain.fi.    3600    IN      MX      20 smtp2.domain.fi.
```

Different fields are explained in Table 1.

Table 1. MX resource record format

Value	Explanation
domain.fi.	The domain that the RR refers to
3600	Time-to-live value of the RR
IN MX	Type of the RR, mail exchanger
10	Priority of the MX RR
smtp1.domain.fi.	Mail exchanger server address

Original DNS specification does not provide any kind of mechanism to proof the authenticity of the response and is vulnerable to man-in-the-middle attacks and DNS cache poisoning attacks.

2.4.2 Simple Mail Transfer Protocol (SMTP)

SMTP was originally specified in RFC 821, published in 1982, but it has since been updated many times. Current up-to-date specification is in RFC 5321 that was published in 2008 (Klensin, 2008). The document specifies a standard protocol for Internet electronic mail transport and mail servers and other message transfer agents use SMTP to send and receive emails.

The original version of SMTP specifies only how SMTP servers can negotiate which each other about the sender, recipient, date, content and other information and report about possible deliv-

ery errors. The protocol is run over TCP and is a plain text protocol without support for authentication or encryption, which makes it vulnerable to eavesdropping and man-in-the-middle attacks if a threat actor is able to monitor the connection or place himself in the middle of the connection.

The protocol has since been extended with various extensions and supports authentication and encryption.

Plain text SMTP transport example can be found in Appendix 1.

2.4.3 STARTTLS

Transport Layer Security (TLS) is a common mechanism for providing confidentiality and authentication to different application protocols using encryption. Its most common application is encrypted HTTP also known as HTTPS.

STARTTLS extension provides this same capability to SMTP protocol (Hoffman, 2002). STARTTLS, also known as opportunistic TLS, provides plain text connections with the opportunity to upgrade to an encrypted connection instead of using a separate TCP port for encrypted communications. TLS is an application independent protocol, so it doesn't have any effect on the core functionality of SMTP after an encrypted channel has been negotiated between the peers.

Opportunistic encryption provides protection from passive monitoring, but the problem with it is that the session is initiated in plain text. If a threat actor can place himself in the middle of the connection to perform a man-in-the-middle attack, the actor can modify the communication and strip the information about STARTTLS support, so that the peers will think that the other end does not support encryption. This attack is known as STRIPTLS attack (Bursztein, 2015).

Example of SMTP transport with STARTTLS extension can be found in Appendix 2.

2.4.4 Domain Name System Security Extensions (DNSSEC)

DNSSEC is an extension to DNS that provides support for cryptographically authenticated responses to domain name lookups to protect from forged responses (NWG, 2005). DNSSEC does not provide confidentiality to DNS queries and responses.

DNSSEC works by adding cryptographic signatures to existing DNS records. These records are stored in DNS name servers alongside with traditional record types like A, AAAA, CNAME, MX, NS and TXT. By checking the associated signature, a DNS resolver can validate that the received DNS responses come from an authoritative name server and wasn't modified while in transit.

This makes it possible for DNS to securely publish different types of security related information in addition to IP address information, such as certificates or their hashes, to be used by other application protocols. DANE is one example of this as it publishes TLS trust anchors in TLSA records.

2.4.5 DNS-based Authentication of Named Entities (DANE)

DANE is a protocol that allow X.509 digital certificates, often used for TLS, to be bound to domain names using DNSSEC (IETF, 2015a). One specific application of DANE is enforcing the usage of STARTTLS and encryption when sending and relaying email messages between SMTP servers. Figure 1 illustrates how DANE works in combination with SMTP.

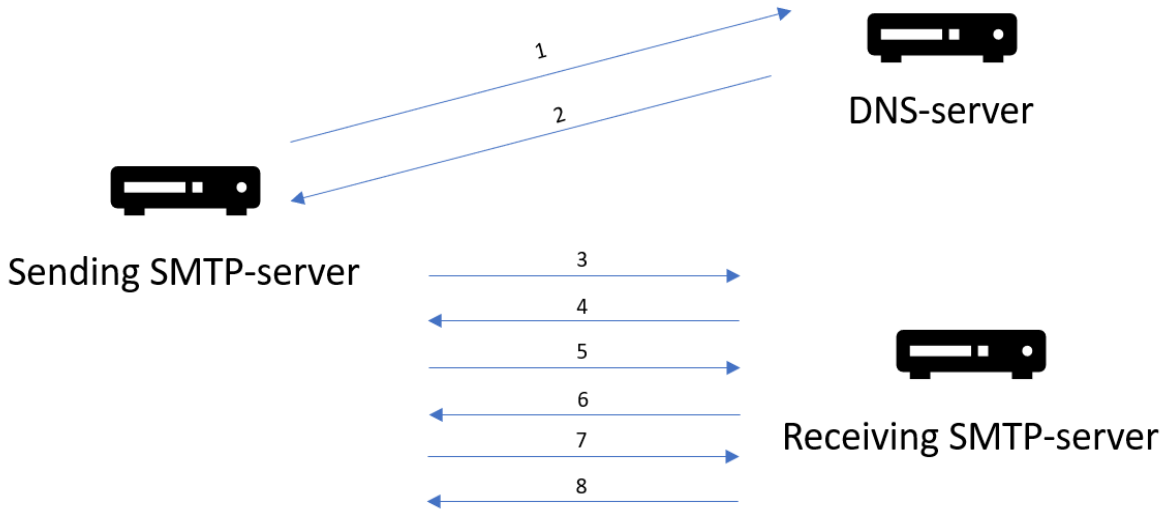

Figure 1. Illustration of how DANE works

1. Sending SMTP server queries TLSA records of the receiving SMTP server through its resolver, which recursively forwards the query to the authoritative DNS server
2. Authoritative DNS server responds with the TLSA records and resolver DNS server uses DNSSEC to validate the response
3. Sending SMTP server establishes a connection to the receiving domains SMTP server as advertised by MX records in DNS to SMTP port (TCP-25)
4. Receiving SMTP server responds with ESTMP banner
5. Sending SMTP server responds with Extended Hello (EHLO)
6. Receiving SMTP server responds with the extensions that it supports, including STARTTLS
7. Sending SMTP server issues a STARTTLS command
8. Receiving SMTP server responds that it is ready to start TLS handshake and the handshake is finished, sending SMTP server compares information of the received certificate to the TLSA information it received in step 2

DANE uses TLSA resource records to publish information about authorized digital certificates in specific services on a specific address. TLSA record has the following format:

```
_25._tcp.smtp1.domain.fi. 3600    TLSA    2 1 1
8D02536C887482BC34FF54E41D2BA659BF85B341A0A20AFADB5813DCFBCF2
86D
```

Different fields are explained in Table 2.

Table 2. DANE resource record format

Value	Explanation
_25	Specifies the TCP or UDP port on the address that this record applies to
_tcp	Specifies the protocol (TCP or UDP) that this record applies to
smtp1.domain.fi.	Specifies the fully qualified domain name (FQDN) of the service
3600	Time-to-live (TTL) of the record, specifies how long the information can be cached on a caching nameserver
TLSA	Specifies the resource record type as TLS Anchor
2	Certificate usage, specifies how to verify the certificate. Possible values are 0-3: - 0 = CA constraint (PKIX-TA), authorizes a root CA or its intermediate certificate. Certificate presented in TLS handshake must have a valid trust chain to a root CA that is already trusted by the client - 1 = Service certificate constraint (PKIX-EE), authorizes a specific certificate published in this record to be used, but in addition the certificate must have a valid trust chain to a trusted root CA - 2 = Trust anchor assertion (DANE-TA), authorizes a certificate that has valid certification path pointing back to the certificate mentioned in this record - 3 = Domain issued certificate (DANE-EE), authorizes a specific self-signed certificate to be used instead of one signed by a CA
1	Selector, specifies which parts of the certificate should be checked. Possible values are 0 or 1: - 0: entire certificate is checked - 1: only public key is checked
1	Matching type, specifies what data is included in the certificate association data. Possible values are 0-2: - 0: entire information is included in the data - 1: SHA-256 data is included in the data - 2: SHA-512 data is included in the data
8D02536C8874…	Certification association data, includes information of the authorized certificate as specified in selector and matching type fields.

Technically DANE could work without DNSSEC, but the added security without DNSSEC would be very limited as the attacker could do a man-in-the-middle attack to both DNS and SMTP sessions and spoof the responses.

Currently DANE has limited support in different SMTP server applications, but many of the widely used service providers already support it.

2.4.6 SMTP MTA Strict Transport Security (MTA-STS)

MTA-STS is an internet standard that instructs an SMTP server that the communications with the other SMTP server must be encrypted and that the domain dame in the certificate must match the domain in the policy (IETF, 2018b). It uses both DNS and HTTPS to publish a policy that tells the sending server that STARTTLS should be used and what to do if an encrypted channel can't be negotiated.

The purpose of MTA-STS is very similar to of DANE, but DANE requires DNSSEC for DNS authentication, while MTA-STS relies on digital certificates signed by trusted certification authorities (CA) (Figure 2). MTA-STS also supports testing mode, so that a domain administrator can start from a reporting only model and proceed to enforcing mode once the configuration is found mature enough. MTA-STS requires the mail server to support TLSv1.2 or higher.

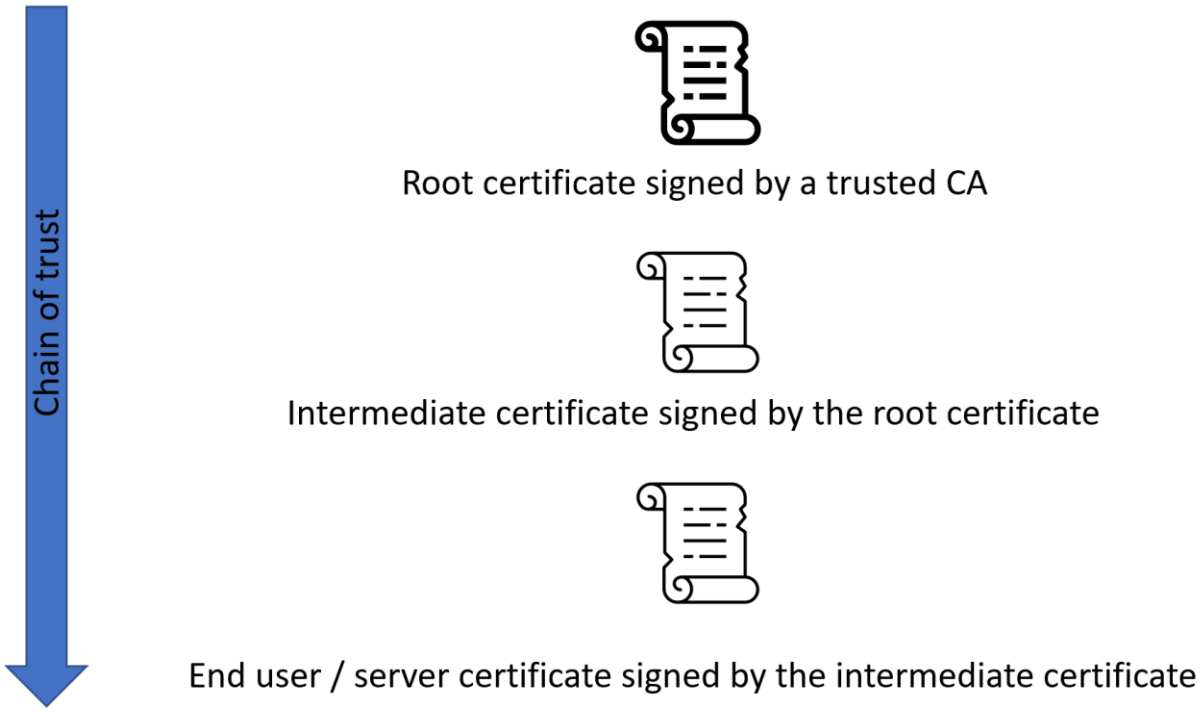

Figure 2. Certificate chain of trust

Illustration of how MTA-STS works can be found on Figure 3.

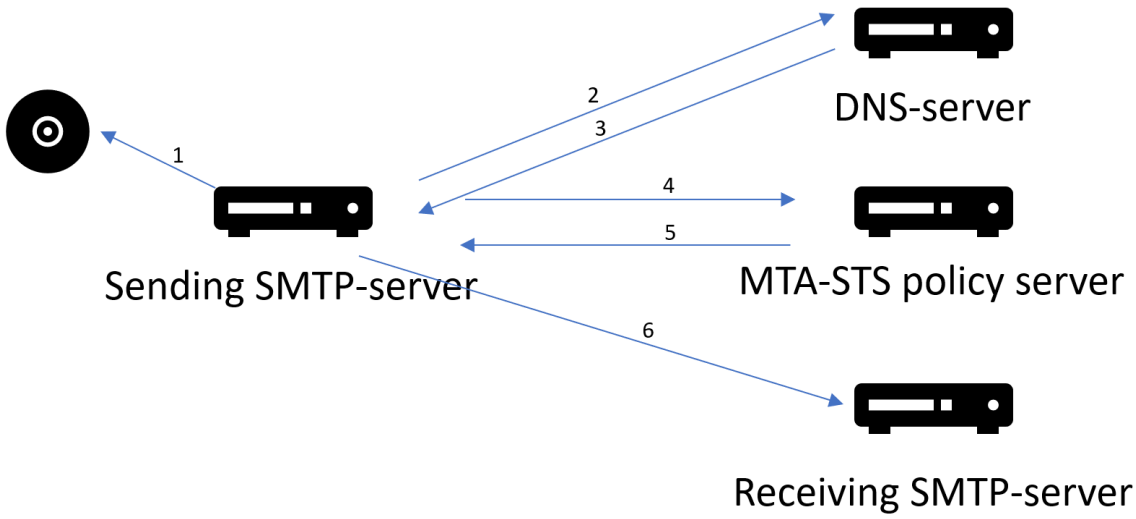

Figure 3. Illustration of how MTA-STS works

1. Sending server checks its policy cache for a policy for the receiving domain
2. Query DNS server for mta-sts and _mta-sts records
3. Authoritative DNS server replies data in the records
4. Sending server tries to download policy file from the server using HTTPS
5. Policy server sends the file
6. Email is sent using enforced TLS (STARTTLS)

MTA-STS queries for two records from the DNS. Mta-sts.domain.fi. should point to the IP address of the policy server, either directly with an A record or via CNAME record:

```
mta-sts.domain.fi.   3600   IN    CNAME   domain.fi.

domain.fi.           3600   IN    A       192.168.30.6
```

_mta-sts.domain.fi. should have a TXT record that contains the version used and a policy ID:

```
_mta-sts.domain.fi. 3600 IN TXT "v=STSv1;
id=20210406194500Z;"
```

Fields are explained in Table 3.

Table 3. _mta-sts TXT record format

Value	Explanation
V=STSv1;	STS version used, currently only STSv1 is supported
id=20210406194500Z;	Id of the published policy. Helps the sending server to determine when the policy has been updated and if the cached policy is still valid

If the policy is not cached by the sending server, or it is expired, a HTTPS request is made to the policy server. RFC defines a standard path for the file:

https://mta-sts.domain.fi./.well-known/mta-sts.txt

And the file has a standard format:

```
version: STSv1
mode: enforce
mx: smtp1.domain.fi
mx: smtp2.domain.fi
max_age: 604800
```

File format is explained in Table 4.

Table 4. MTA-STS policy file format

Value	Explanation
version	STS version used, currently only STSv1 is supported
mode	Indicates expected behavior when a policy validation failure is detected. Valid values are: - enforce: sending server will not deliver the mail to receiving server - testing: mail will be delivered but a report will be sent if TLS-RPT is configured - none: policy is ignored, used for removing MTA-STS
mx	One or more patterns matching the MX records of a receiving domain
max age	Maximum lifetime of the policy in seconds. Mail servers will cache policies for the time specified here

Currently support for MTA-STS is limited, but major cloud services providers like Google's gmail and Microsoft's O365 have recently started to support it.

2.4.7 TLS Reporting (TLS-RPT)

DANE and MTA-STS only provide a mechanism to enforce encryption and authenticate the receiving SMTP server, but they don't include functionality for reporting errors to the receiving server. To overcome this problem a new protocol called TLS-RPT was defined by IETF (IETF, 2018b). It allows a domain to advertise a destination for the sending email servers to report the success or failure of encryption in SMTP connections.

TLS-RPT is configured by publishing a TXT resource record in DNS at _smtp._tls.domain.fi:

```
_smtp._tls.domain.fi.  3600    IN      TXT     "v=TLSRPTv1;
rua=mailto:xxxxxxx@tls.eu.dmarcian.com"
```

Format of the record is explained in Table 5.

Table 5. TLS-RPT format explained

Value	Explanation
v=TLSRPTv1	Version indicator, currently v1 is the only supported version
rua=mailto:r1vm3aek@tls.eu.dmarcian.com	Defines the address to which the reports are delivered. Delivery protocol can be either mailto or http

After a TLS-RPT record is published in DNS, sending mail servers that support it will send a daily aggregate report to the specified address about the messages that they've sent. The reports are compressed JSON files that are difficult to read for humans, but there are services that convert the information into more human readable format. Example report from a service provider can be seen in Figure 4.

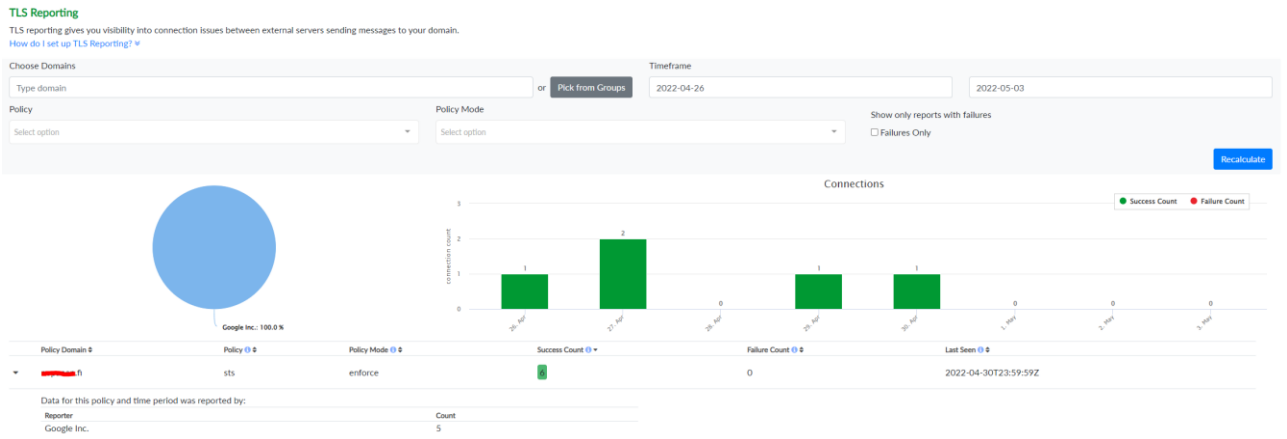

Figure 4. TLS report from dmarcian

TLS-RPT supports the reporting of several different errors related to TLS negotiation failures (Table 6), MTA-STS related failures (Table 7) and DNS related failures (Table 8).

Table 6. TLS negotation failures

Reason	Explanation
starttls-not-supported	The receiving MTA does not support the STARTTLS command
certificate-host-mismatch	The receiving MTA's certificate is not matching the hostname
certificate-not-trusted	The sender does not trust the certificate, which the receiving MTA supplied
certificate-expired	The receiving MTA's certificate is expired
validation-failure	Other general validation failures than the ones mentioned above

Table 7. MTA-STS related failures

Reason	Explanation
sts-policy-fetch-error	The sender could not fetch the MTA-STS policy over HTTPS
sts-policy-invalid	indicates a syntax error in the policy, prevent-ing the validation of the MTA-STS policy
sts-webpki-invalid	failure to fetch the MTA-STS policy due to PKI validation issues

Table 8 DNS related failures

Reason	Explanation
tlsa-invalid	indicates a TLSA record validation error
dnssec-invalid	shows the inability of the recursive resolver to return a valid record
dane-required	suggests that the sending domain requires DANE TLSA records of the destination domain (MX hosts), but it could not find any DNSSEC-validated TLSA records

2.5 Information gathering

The starting point of this thesis work was to obtain a complete list of fi-domains from the Finnish domain authority Traficom for research purposes, and then gather all the required information utilizing readily available open-source tools as much as possible. Several custom scripts were required to automate the information collection and analysis.

The complete listing of *.fi* domains (.fi zone file) is considered to be confidential and is not publicly available, but all other information was obtained from publicly available DNS and HTTP servers and is available to everyone who knows the domain name.

Authoritative DNS servers as pointed out in the zone file were queried for the required resource records and recursive resolving was used in case the records referenced to external zones. Required information from the responses was stored in different files for analysis. Servers that did not reply within the timeout threshold labeled and separated in the results.

2.5.1 Setup for information gathering

An OpenBSD 7.0 server was used as the platform. Server was connected to the Internet using DNA Ltd. as the internet service provider and the server had fixed IPv4 and IPv6 addresses. Unbound DNS validating resolver 1.13.2 was used as a DNS resolver with DNSSEC validation enabled so that all DNSSEC signed responses were checked and invalid ones discarded.

Dig 9.10.8-P1 DNS lookup utility was used to perform individual DNS queries and curl 7.79.0 was used to perform required HTTPS request.

2.5.2 DANE information gathering

Checking for DANE was done in four steps:

- Check if the domain has DNSSEC configured
- Check if the domain has MX records
- Check if the SMTP server's that MX records points to have TLSA records
- Check if there are TLSA records for none, some, or all SMTP servers

Because DANE requires DNSSEC to work (IETF, 2015a), first step was to include only those fi-domains that had DNSSEC configured. This was done by collecting domain names with DS resource records (RR, records from here after) set in the *.fi* root zone file. DS record includes a hash of the domains DNSKEY and is submitted by the domain operator to the parent zone operator to be published in the DNS. It allows the transfer of trust from parent zone (fi) to a child zone (individual fi-domain) and thus is a good indicator of a domain with DNSSEC enabled.

An example of DNSSEC enabled domain in the zone file:

```
domain.fi.      IN      NS      ns1.domain.fi.
                IN      NS      ns2.domain.fi.
domain.fi.      IN      DS      1076    10      2
c80bc621faeeeb941…3119548f9a2383918b68d2815f5
domain.fi.      IN      DS      33147   10      2
c930bedc4b4c15f9a…c8e8f10e7aa03537cc36beb999f
```

Next step was to check if the DNSSEC enabled domains have MX records published, which indicates that they want to receive email. This was done using dig to query recursive resolver:

```
dig +noall +answer -t MX domain.fi.
```

Responses were stored in text files to be used in the next step. As a DNSSEC validating resolver was used, domains that did not pass the validity test were abandoned already at this stage.

If the response was empty, the domain was considered not to have any MX records in place. If the domain replied with one or more MX record, the next step was to check for NULL MX records. Domain operator can publish a null MX to prevent email delivery to the domain (IETF, 2015b).

Example NULL MX record:

```
domain.fi.              294     IN      MX      0 .
```

In these cases, no further checks were done.

Example of a domain with two MX records:

```
domain.fi.    3600    IN      MX      0 smtp1.domain.fi.
domain.fi.    3600    IN      MX      10 smtp2.domain.fi.
```

If the domain replied with one or more MX record, the next step was to check if there were matching TLSA records for the SMTP servers. DNS query was done with adding string "_25._tcp." In front of the hostname included in the MX record.

Dig query to check for TLSA records on MX server:

```
dig +noall +answer -t TLSA _25._tcp.smtp1.domain.fi.
```

Responses were stored in text files for later comparison.

Example reply containing TLSA records:

```
_25._tcp.smtp1.domain.fi. 3600   IN      TLSA    2 1 1
276FE8A8C4EC7611565BF9FCE6DCACE9BE320C1B5BEA27596B2204071
ED04F10

_25._tcp.smtp1.domain.fi. 3600   IN      TLSA    2 1 1
60B87575447DCBA2A36B7D11AC09FB24A9DB406FEE12D2CC901805176
16E8A18

_25._tcp.smtp1.domain.fi. 3600   IN      TLSA    2 1 1
8D02536C887482BC34FF54E41D2BA659BF85B341A0A20AFADB5813DCF
BCF286D

_25._tcp.smtp1.domain.fi. 3600   IN      TLSA    2 1 1
BD936E72B212EF6F773102C6B77D38F94297322EFC25396BC3279422E
0C89270

_25._tcp.smtp1.domain.fi. 3600   IN      TLSA    2 1 1
E5545E211347241891C554A03934CDE9B749664A59D26D615FE58F779
90F2D03
```

A domain could have matching TLSA records for none, some, or all of it MX records, so the final step is to check if the number of hosts with TLSA record match the number of MX records and label the result as none, partial or all.

2.5.3 MTA-STS information gathering

Checking for MTA-STS was done in three steps:

- Check if the domain has hostname mta-sts.domain.fi. in the DNS
- Check if there is a valid MTA-STS configuration file published in URL https://mta.sts.domain.fi/.well-known/mta-sts.txt
- Check if the domain has a valid TXT record for _mta-sts.domain.fi

MTA-STS information gathering started by doing a DNS lookup for all fi-domains with mta-sts added to the hostname:

```
dig +noall +answer mta-sts.domain.fi.
```

Responses were stored in text files to be used in the next steps. Domains without a record for mta-sts were discarded at this phase.

Example of a valid answer:

```
mta-sts.domain.fi.    3600   IN    CNAME    domain.fi.
domain.fi.            3600   IN    A        192.168.100.100
```

Next step was to try to download `mta-sts.txt` file from the server using curl:

```
curl --connect-timeout 7 --max-time 20 --silent -fo do-
main.fi. https://mta-sts.domain.fi./.well-known/mta-
sts.txt
```

Successfully downloaded files were saved and analyzed if they had a valid `mta-sts.txt` file format. All possible errors were recorded based on exit status of curl for further analysis.

Example of a valid `mta-sts.txt` file:

```
version: STSv1
mode: enforce
mx: smtp1.domain.fi
mx: smtp2.domain.fi
max_age: 604800
```

Because MTA-STS also needs a valid TXT record under _mta-sts.domain.fi., last step was to check for that:

```
dig +noall +answer -t txt _mta-sts.domain.fi.
```

Responses were stored I text files to be compared with the results done in the previous steps.

Example of a valid _mta-sts TXT record:

```
_mta-sts.domain.fi.   3600   IN   TXT   "v=STSv1;
id=20210406194500Z;"
```

As MTA-STS needs both a `mta-sts.txt` file and _mta-sts TXT record in order to work, the final step was to combine the results and check which domains had all the required information published.

2.5.4 TLS-RPT information gathering

Checking for TLS-RPT was performed by doing a single DNS query to each domain that had DANE, MTA-STS, or both configured:

```
dig +noall +answer -t txt _smtp._tls.domain.fi.
```

Responses were stored in text files and analyzed to see if it included a valid _smtp._tls TXT record:

```
_smtp._tls.domain.fi.  3600   IN   TXT   "v=TLSRPTv1;
rua=mailto:xxxxxxx@tls.report.service.com"
```

2.6 Study ethics

The study is conducted by performing DNS and HTTPS requests over the public Internet. In both cases there is a TCP or UDP connection established from the information gathering server to the target server and it can be seen for example on the firewall or application logs. On the other hand, information queried for the study is by definition public and required to deliver email to the domain and thus many SMTP servers sending mail to the domain will perform the same queries.

Server performing the queries had RIPE whois-records and PTR records for both IPv4 and IPv6 addresses that can easily be connected to the author of the study.

Information about the full *.fi* zone could be seen as valuable to someone performing information gathering, but most of the information can be fetched from Traficom's public API's, Certificate Transparency lists and with brute force enumeration.

Results of the study indicate only statistics and no information about individual domains are revealed.

3 Threats to the confidentiality of SMTP connections

As email is a highly distributed system with many different actors in different roles as shown in Figure 5, there are many possibilities for a threat actor to compromise the confidentiality of communications. A malware in the user's device can leak information about the emails to a threat actor, malicious system administrator could access users mails at the server's mail queue or mailboxes in the mail server, or someone being able to eavesdrop on the traffic can monitor email traffic between SMTP servers.

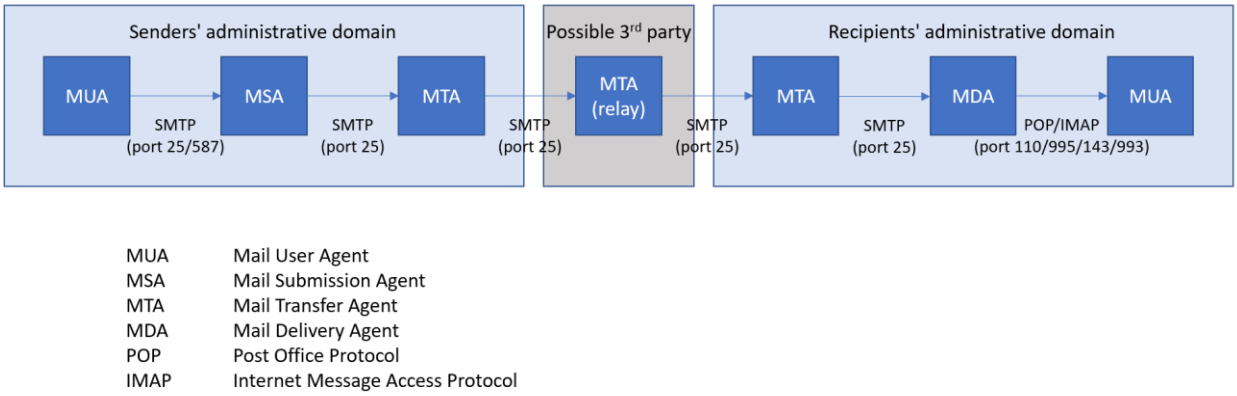

Figure 5. Illustration of email delivery path

Some example threats are explained in more detail in this chapter.

3.1 Threats against MUA

The number of different computer viruses and malicious software has exploded since the beginning of home computing. Modern malicious software includes functionalities that are especially targeting the users email client and tries to send data from the client to the attacker, including for example address books and actual emails.

Older versions of mail user agents used plain text SMTP when sending email messages and POP/IMAP when receiving messages. These are nowadays mostly replaced with TLS encrypted versions.

Enforcing SMTP encryption in transit with the techniques presented in this study does not prevent these kinds of breaches. Most efficient controls include keeping your operating system and applications updated, having up-to-date antivirus software installed and, in some cases, encrypting confidential emails with separate encryption application.

3.2 Threats against MSA and MTA

Sending, relaying, and receiving SMTP servers usually see individual emails in plaintext as they need to access the headers to be able to deliver them to the correct recipient. Message content is not accessible only if it is encrypted at the application layer, using for example S/MIME (NWG, 2009) or PGP (NWG, 2007). This makes it possible for a malicious server administrator or malicious software installed on the server to breach confidentiality and access the content of emails. Enforcing encryption between the SMTP servers does not prevent these kinds of attacks.

3.3 Intercepting SMTP traffic on transit

SMTP to SMTP server communication is considered to be the most vulnerable component of the email system, according to the JRC 2015 report. Originally SMTP protocol was built under the assumption that the servers can trust each other, and no additional security features were built into it. When a sending SMTP server contacts a receiving SMTP server as instructed in the domains MX records to deliver the message, there is an implicit assumption of trust, and no authentication or encryption is performed.

Three example types of threats posed to server communications are passive eavesdropping, on-the-fly tampering, and active interception of SMTP communications.

Because SMTP connections between different organizations usually go through the public Internet, depending on the location of the mail servers' the traffic often goes through different Internet service providers and can cross multiple national borders. Many countries have legislation that allows lawful interception and mass surveillance of cross-border traffic and because SMTP has traditionally been plain text protocol transmitted over TCP connections, even massive eavesdropping of email communications in such cases is trivial. In addition to governmental surveillance, this kind of surveillance can also be found in the context of industrial espionage or criminal activities (JRC, 2015).

If a threat actor can perform a man-in-the-middle attack to the traffic, for example by using DNS cache poisoning, it has even more possibilities: It can try to downgrade the connection to plaintext by performing STRIPTLS attack using on-the-fly tampering, even if the mail servers would otherwise support encryption using STARTTLS, or it could mimic the receiving SMTP server and sup-

port encryption using STARTTLS and intercept the emails before relaying the to the proper receiving server (active interception). In this case, integrity of the emails would also be in danger as the attacker could modify the headers or content of the messages without the recipients knowing it.

Border Gateway Protocol (BGP) is the core routing protocol of Internet. It enables different autonomous systems to find a path between each other via other autonomous systems. In some cases, advanced threat actors can try to alter the route of the traffic using technique called BGP hijacking (Cloudflare, 2022), which allows them to route traffic through infrastructure controlled by them, even if it normally would go through a different route. This would allow the threat actor for example to modify DNS responses to include altered MX records that points to SMTP servers controlled by the attackers, or intercept SMTP traffic and perform a STRIPTLS attack as mentioned above.

One recent practical example of rerouting and possibly intercepting Internet traffic comes from occupied Ukrainian territory in May 2022 when Russia rerouted internet service in Kherson, Ukraine through Russia's network instead of Ukrainian telecommunications infrastructure (The Record, 2022)

DANE and MTA-STS, when deployed correctly and together, are very efficient security controls to prevent SMTP eavesdropping in situations when there is a risk that a threat actor might have access to the traffic.

4 Comparison points

4.1 Overview of policies and recommendations

Research about published policies and recommendation regarding DANE and MTA-STS was made. Criteria for the selection of sources was that the policy or recommendation must be publicly available, published by a credible source (government authority, well known company, or organization) and is published or updated recently. Summary of the recommendations is presented in Table 9.

The Central Digital and Data Office United Kingdom (gov.uk) has published a guidance on "Set up government email services securely" (UK Government, 2021). The guidance says, that "Government email administrators should follow this guidance to implement encryption and anti-spoofing" and describes different techniques to do that. The first recommendation in the guidance is to set up encryption of emails in transit by enabling TLS on the organizations mail server and enforcing it using MTA-STS. The guide specifically says that once successfully configured, MTA-STS mode should be set to enforce. The guide doesn't mention DANE.

National Cyber Security Centre of UK (NCSC-UK) has published its own guidelines that are fully in line with the government's guidance (NCSC, 2019).

Australian Cyber Security Centre (ACSC) has published a similar guideline under the topic "Guidelines for email" (ACSC, 2022). The guideline recommends using opportunistic TLS encryption (STARTTLS) but reminds about the risk of STRIPTLS attack and recommend implementing MTA-STS to reduce it. DANE is not mentioned.

Canadian Centre for Cyber Security has a publication titled "Implementation guidance: email domain protection (ITSP.40.065 v1.1)" (Canadian Centre for Cyber Security, 2021). The guideline makes three recommendations regarding email transport encryption:

1. STARTTLS
2. DNS-Based authentication of named entities (DANE)
3. MTA Strict transport security (MTA-STS)

Dutch Forum Standaardisatie publishes a list of mandatory standards that must be implemented by government bodies (Forum Standaardisatie, 2022). The list includes STARTTLS and DANE to secure connection between mail servers. MTA-STS is not mentioned.

Table 9. Summary of DANE and MTA-STS recommendations

Source	DANE	MTA-STS
gov.uk	no	yes
NCSC-UK	no	yes
ACSC	no	yes
Canadian Center for Cyber Security	yes	yes
Forum Standaardisatie	yes	no

It can be summarized that surprisingly MTA-STS is mentioned more often than DANE, even though it is a newer protocol. One reason explaining this might be that DANE requires DNSSEC to work and implementing that is sometimes considered to be a complex process. With MTA-STS, only a certificate signed by a trusted CA is needed and those are nowadays easily and freely available for example from Let's Encrypt.

4.2 Statistics of implementation rate

There are some statistics available on the usage of email security technologies covered in this study, but there aren't many ccTLD wide statistics publicly available, that could work as a direct comparison point to *.fi* domains. Some of the statistics that are available are explored in this chapter.

4.2.1 Google

Google publishes up-to-date statistics about email encryption in transit to and from their mail servers (Google, 2022). These statistics include only usage of STARTTLS and does not include usage of DANE or MTA-STS.

The statistics in Figure 6 and Figure 7 show that on average close to 90% of the incoming and outgoing mails are encrypted.

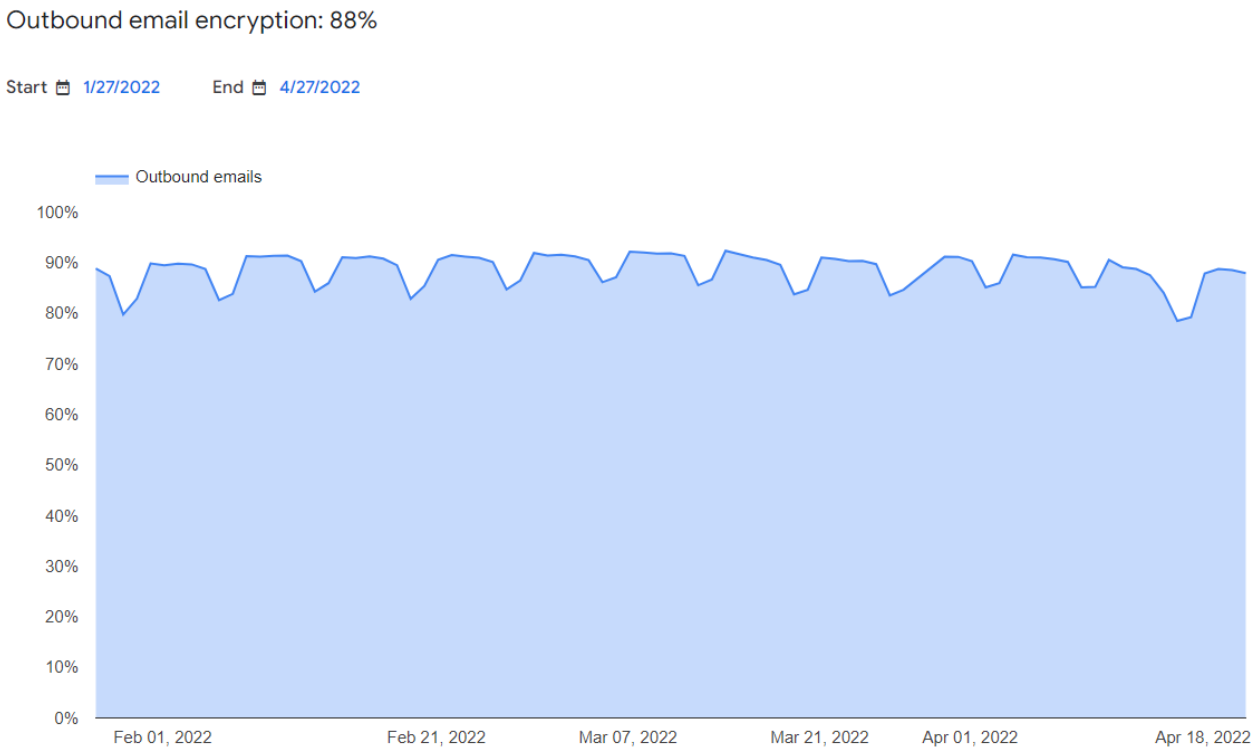

Figure 6. The percentage of TLS encrypted emails outbound from Google

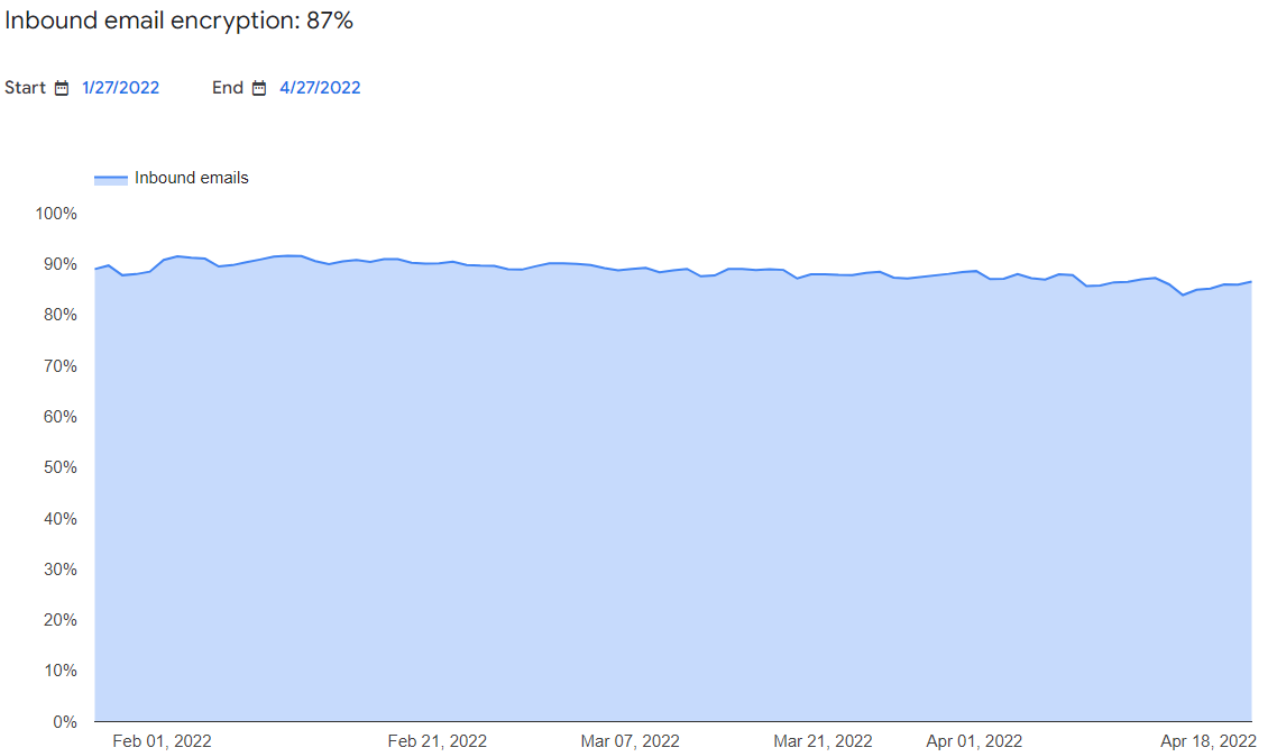

Figure 7. The percentage of TLS encrypted emails inbound to Google

4.2.2 DNSSEC-Tools

DNSSEC-Tools project collects daily statistics about DNSSEC and DANE deployments (DNSSEC-tools, 2022). The latest statistics are shown in Table 10. Statistics are collected on a global scope and the project has access to some individual TLD zone files, but it doesn't publish statistics about individual zones. Statistics about MTA-STS are not collected.

Table 10. The latest numbers from DNSSEC-tools project

Last updated	2022-04-26 04:50 -0700
Total number of DS resource record sets	18,343,379
Total number of working DNSKEYs	18,160,492
Total DANE protected SMTP domains	3,188,686

Figure 8 shows the growth of observed DS record (i.e., the number of signed zones) sets over time.

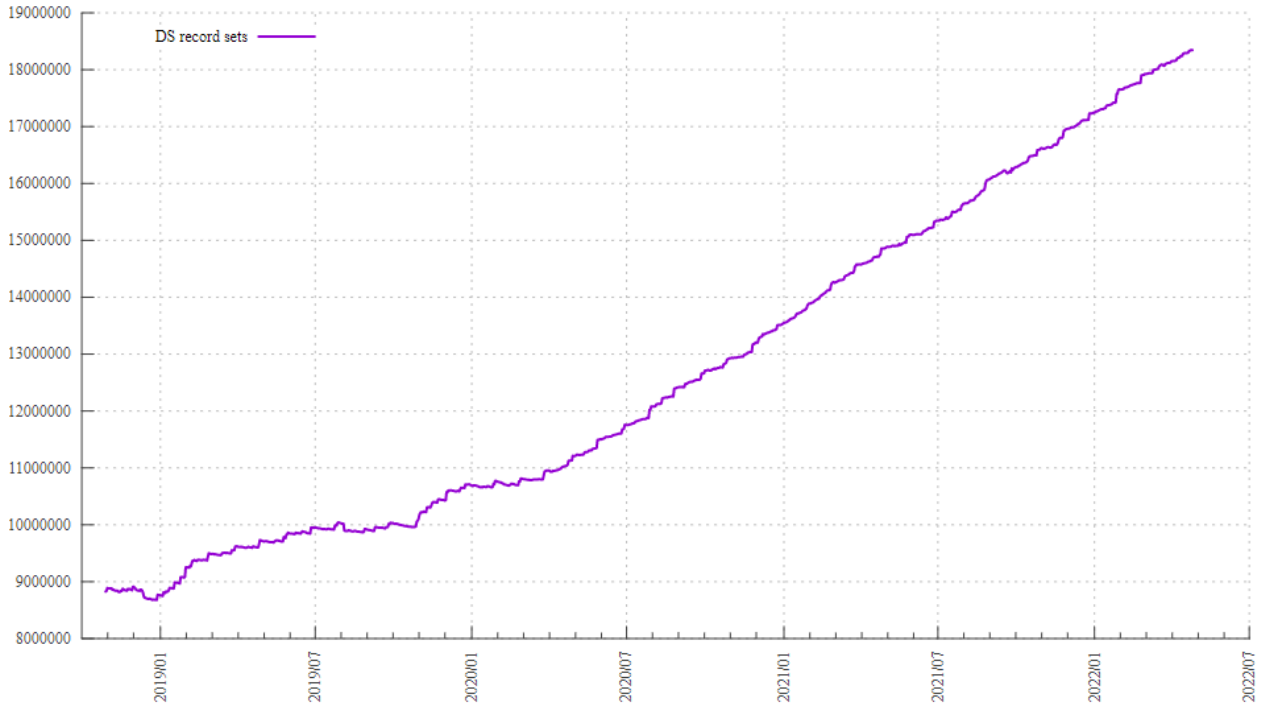

Figure 8. Growth of observer DS record sets over time

Figure 9 depicts the number of domains that have deployed DANE/SMTP. Specifically, their zone is signed, and their MX records all point to hosts that have DANE TLSA records.

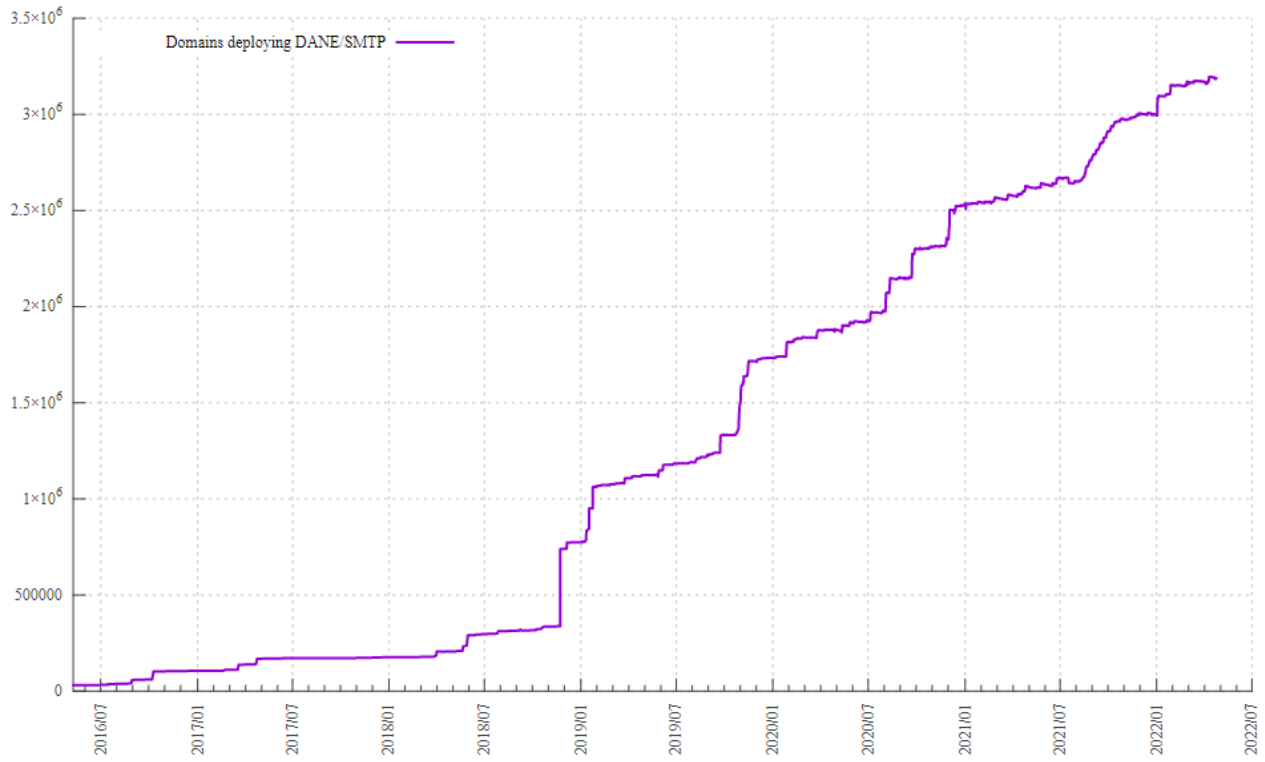

Figure 9. The number of domains that have deployed DANE

Many of the domains in the previous graph outsource or aggregate their mail servers such that the MX records point to out-of-domain mail servers (i.e., externally hosted mail providers). Figure 10 depicts the number of DNS mail provider zones that have DANE records deployed for them. Out of more than 3 000 000 domains with DANE configured, about 8000 domains actually host SMTP servers inside their own domain name.

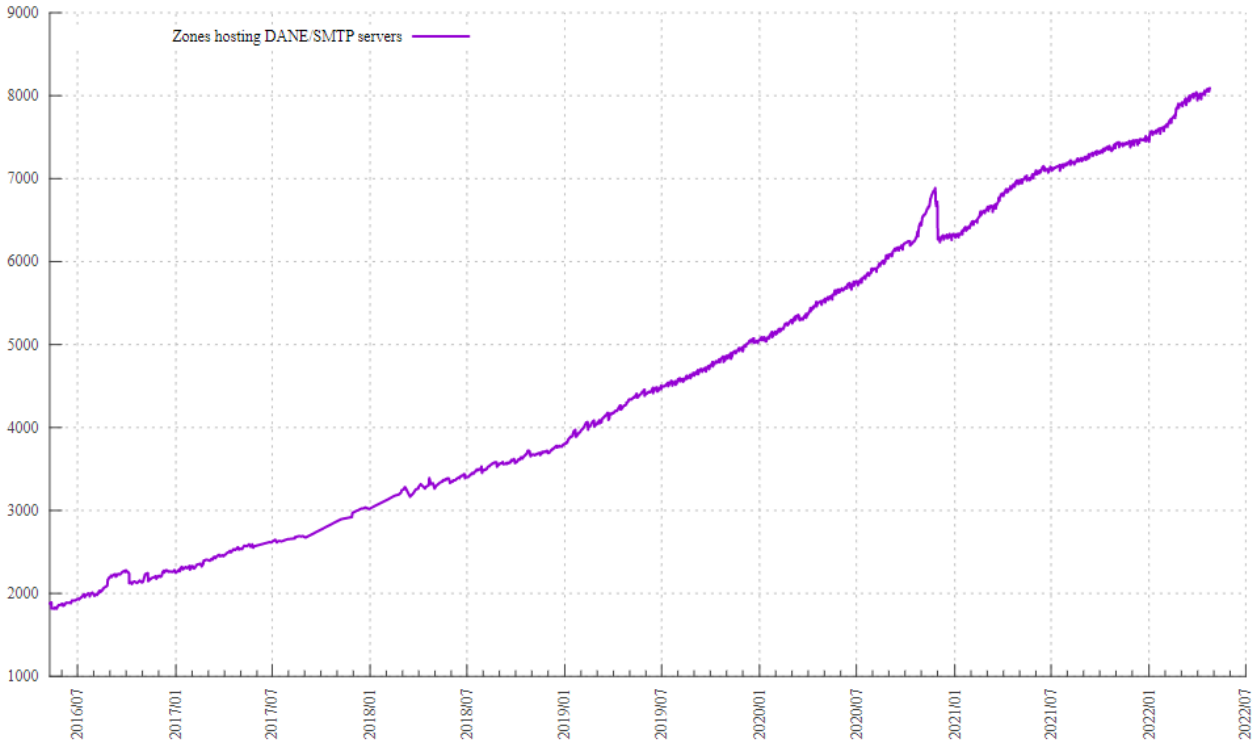

Figure 10. The number of zones hosting DANE protected mail servers

4.2.3 SIDN Labs

SIDN is the national authority for *.nl* ccTLD and SIDN Labs is their research team, whose aim is to contribute to ongoing improvement of the internet infrastructure's trustworthiness through technical applied research. They collect regular statistics on the DANE usage from the *.nl* zone (SIDN Labs, 2022). No statistics about MTA-STS are collected.

Figure 11 shows that as of 20.4.2022 878 761 or 14.08% out of 6 243 034 domains had DANE configured.

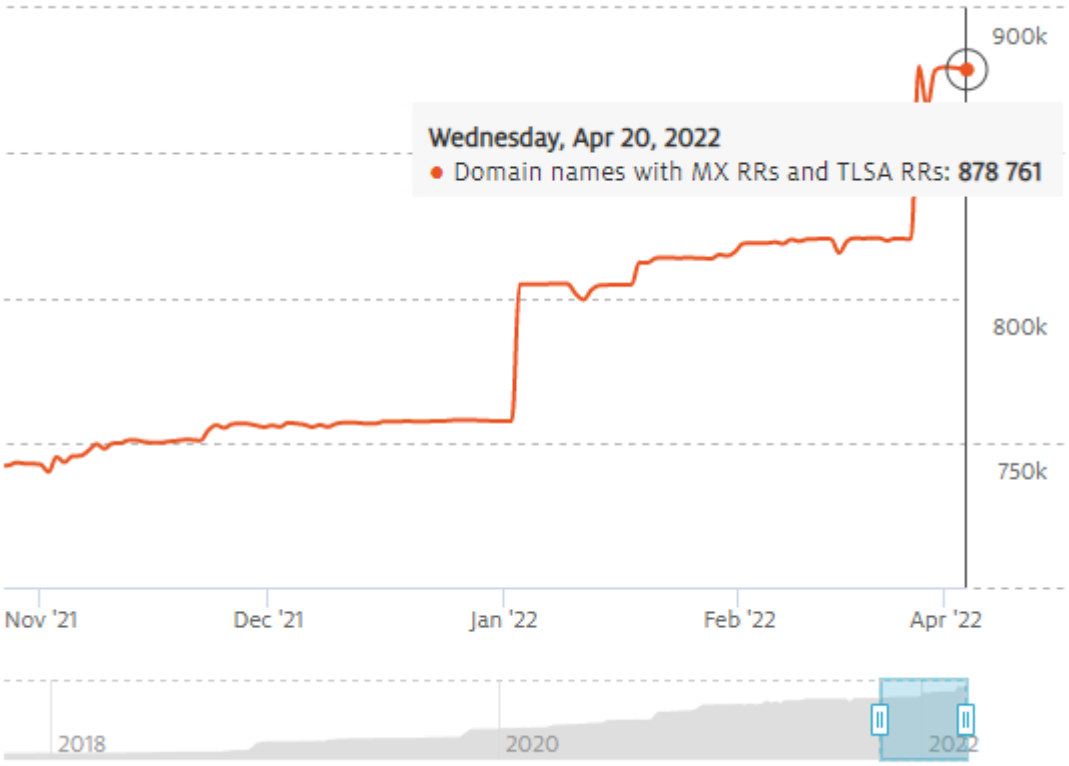

Figure 11. DANE protected domains in *.nl* zone

A common characteristic for statistics collected by both DNSSEC-Tools and SIDN Labs is rapid and occasionally even vertical growth. According to SIDN, the explanation for that is that major email service providers like one.com and TransIP have bulk-enabled DANE for the mail domains that they manage (SIDN, 2019). As an example, bulk operation performed by one.com can be seen in Figure 9 in November 2018. The relatively low number of zones hosting DANE protected mail servers compared to the high number of domains that have deployed DANE for SMTP confirms this conclusion.

According to the same source, high adoption rates in Netherland can be explained by the country's lead taken on DNSSEC signing, which is a prerequisite for the use of DANE. Both DNSSEC and DANE are in the 'use-or-explain' list of the Forum Standaardisatie, making it more or less mandatory for government organizations to deploy them.

4.2.4 Microsoft O365

Microsoft doesn’t publish detailed or updated statistics, but they announced in a blog article dated 2.2.2022 that “Currently, we successfully validate connections to over 35K MTA-STS-protected domains, and this number is growing every month” (Microsoft, 2022).

5 Current status

The starting point for the data collection was a complete *.fi* zone file dated 4th of April 2022, obtained from Traficom. Data was collected using methods described in chapter 2.5 during week 15/2022. The number of domains in the zone file was 524728.

Scripts used to collect the data were developed in agile method in iterations and several test runs were made to reduce the possibility of errors. Some corner cases were confirmed manually, and random manual checks were done to the data. Final data collection was done once as there were time constraints that prevented a longer-term observation.

5.1 DANE

As DANE requires DNSSEC to be configured for it to work, the first step was to extract the domains that had DNSSEC configured from the full zone. This was done by checking which domains had a DS record set. The number of such domains was 18127. Statistics gathered from those domains are shown in Table 11.

Table 11. DANE-statistics for *.fi* zone

Explanation	Count	Percentage	
Number of *.fi*-domains	**524728**		
dnssec configured *.fi* domains	18127	3.45%	of fi-domains
dnssec validation failed	62	0.34%	of dnssec configured fi domains
MX record DNS query timed out	8	0.04%	of dnssec configured fi domains
no MX records configured	3913	21.59%	of dnssec configured fi domains
MX records exists	14144	78.03%	of dnssec configured fi domains
MX record is NULL	2923	20.67%	of MX records exists
MX record is not NULL	11221	79.33%	of MX records exists
TLSA query timeout	3	0.03%	of MX record is not NULL
no TLSA records configured	6755	60.20%	of MX record is not NULL
TLSA records configured	4463	39.77%	of MX record is not NULL
TLSA record for all MX records	**4441**	99.51%	of TLSA records configured
TLSA record for some MX record	**13**	0.29%	of TLSA records configured
invalid TLSA record	9	0.20%	of TLSA records configured

Figure 12 shows that 18127 or 3.45% of fi-domains have DNSSEC configured.

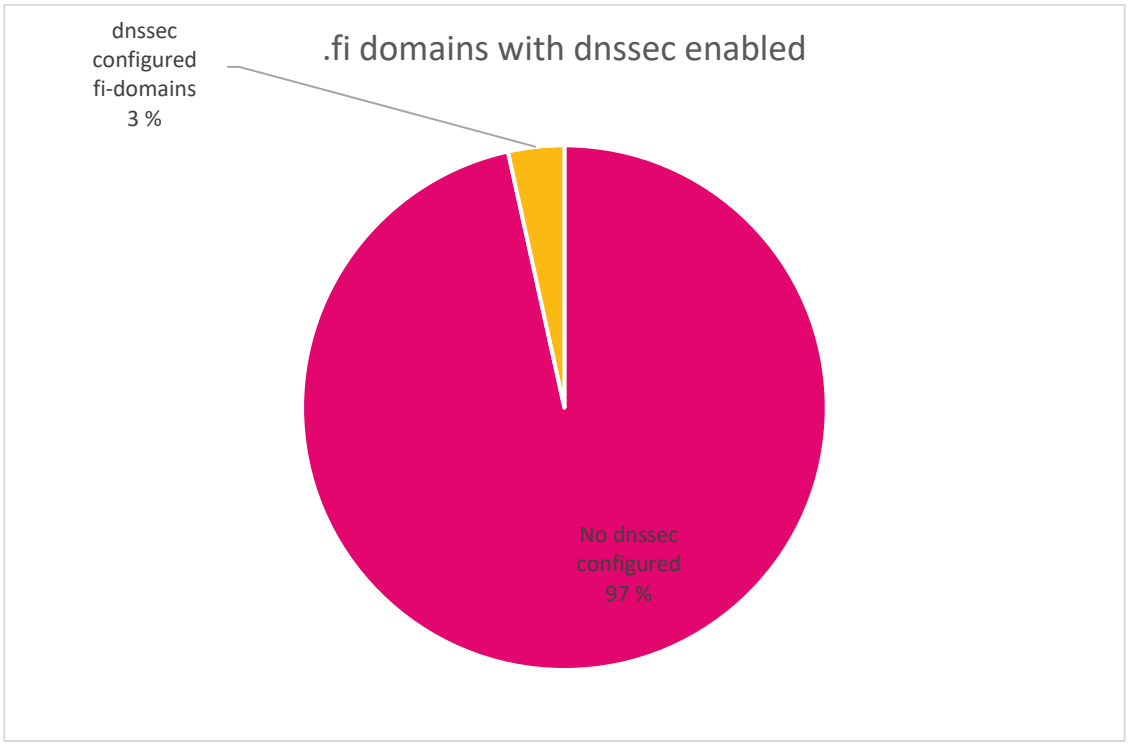

Figure 12. *.fi* domains with dnssec enabled

Out of those domains, 14144 or 78.03% have a MX record configured that is required to route emails to the domain (Figure 13).

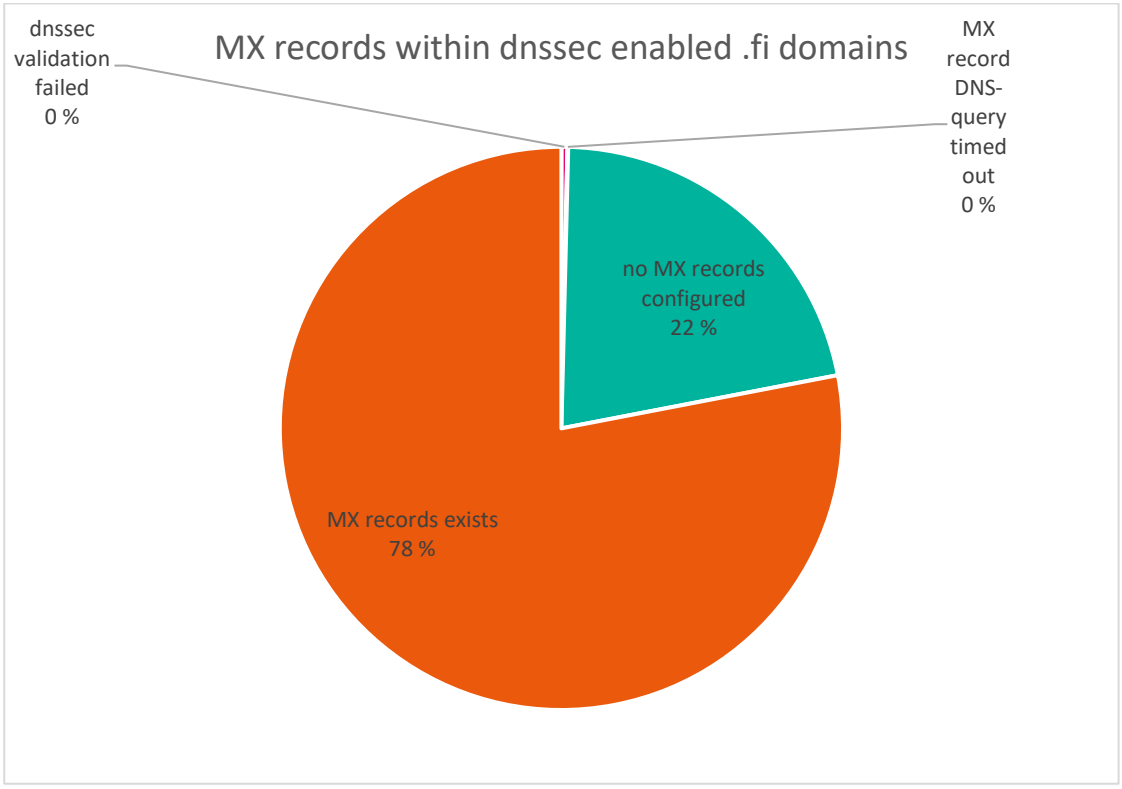

Figure 13. MX records within dnssec enabled *.fi* domains

Domains with NULL MX record imply that the domain owner doesn't want to receive emails and thus DANE is not relevant, so 11221 domains with not-NULL MX records remain (Figure 14).

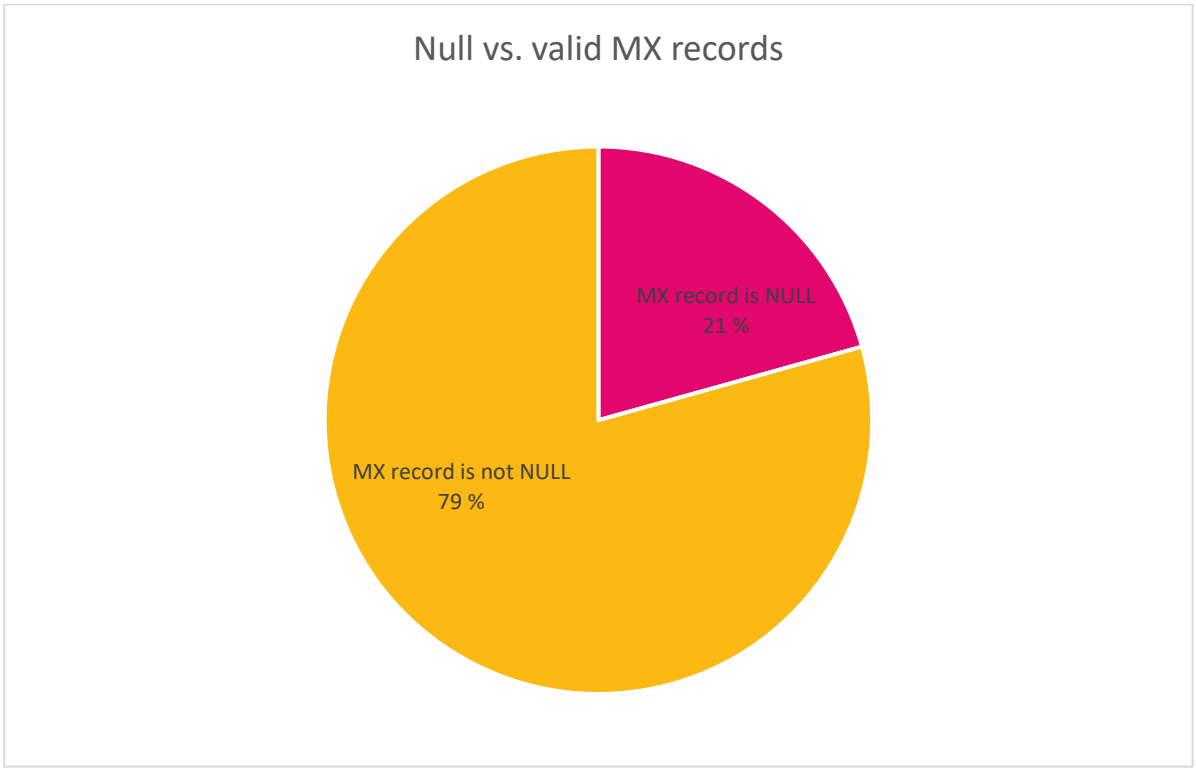

Figure 14. Null vs. valid MX records

Out of domains with not-NULL MX record, 4463 have at least some TLSA records configured for MX records (Figure 15).

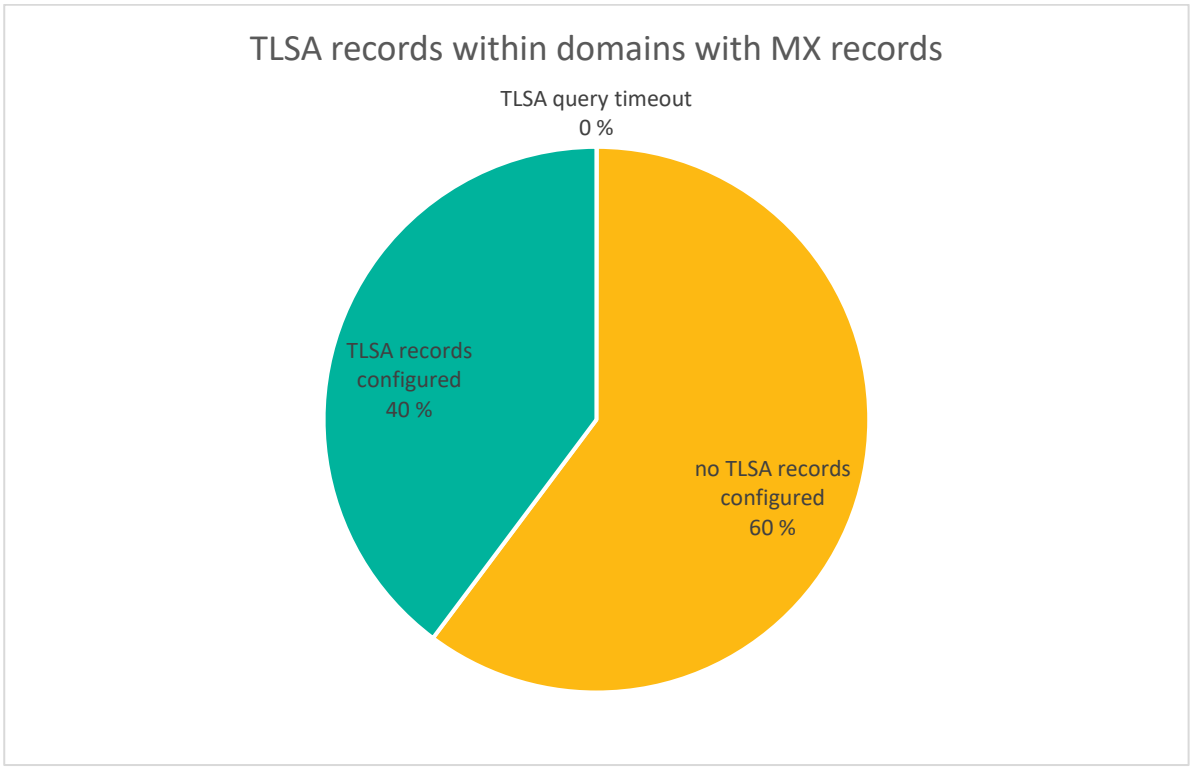

Figure 15. TLSA records within domains with MX records

From domains with TLSA configured for at least some MX records, 4441 have them for all MX records, 13 have them for some but not all MX records and 9 have invalid TLSA records (Figure 16).

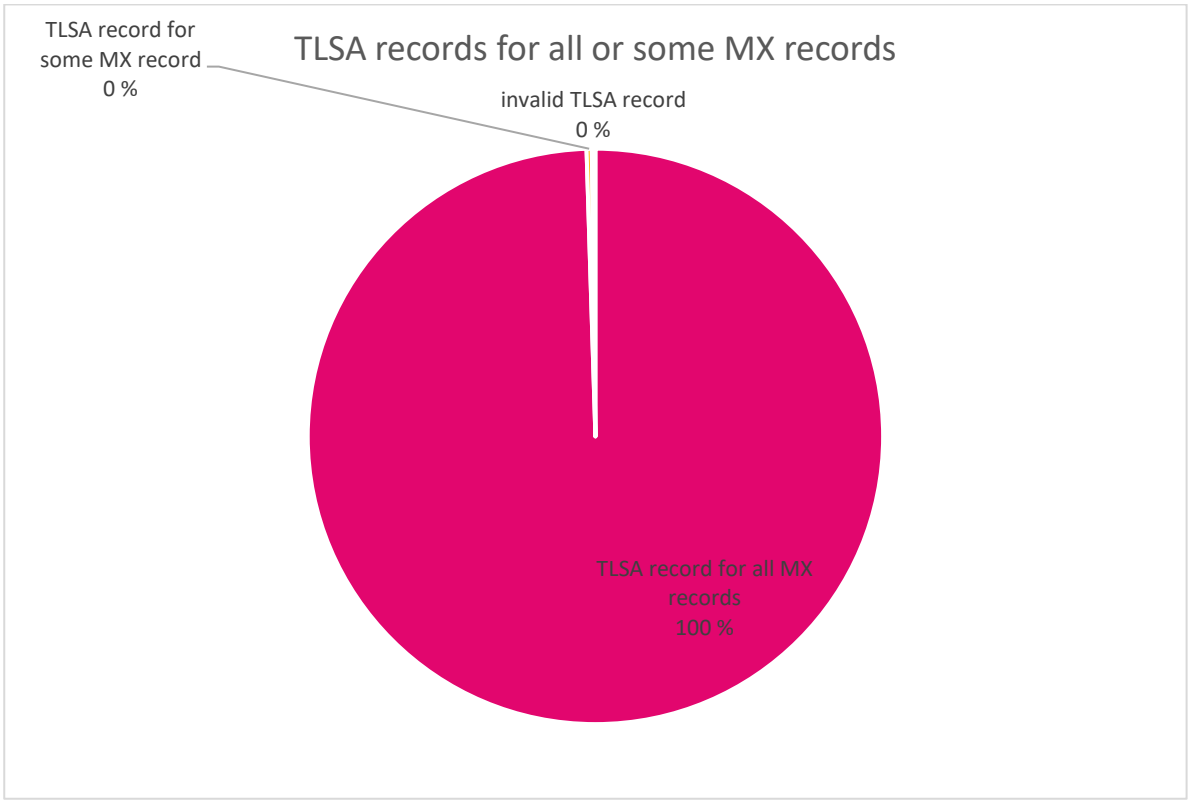

Figure 16. TLSA records for all or some MX records

The majority of the DANE-enabled MX-servers are hosted by service providers as there are only 164 unique MX servers within the MX records protected by TLSA record. The most popular MX server is present in 1255 out of the 4454 domains with full or partial TLSA information.

A key finding regarding DANE implementations in the *.fi* zone is that it is still very limited as can be seen in Table 12.

Table 12. DANE implementations in *.fi* zone

Explanation	Count	Percentage
fi-domains with TLSA records for all or some MX records	4454	
Out of dnssec enabled fi-domains		24.57%
Out of all fi-domains		0.85%

Four key findings can be summed up from the above figures:

1. Implementing DANE to protect receiving SMTP servers in *.fi* zone is very limited as only less than one percent of the domains have it configured to at least some of the SMTP servers.
2. Organizations and individuals who have implemented DNSSEC to protect their DNS are almost 29 times more likely to implement DANE compared to whole *.fi* zone.
3. Low number of DNSSEC enabled domains is limiting the possibility to implement DANE
4. One service provider is hosting mail servers for the majority of the DANE enabled domains

5.2 MTA-STS

Data collection for MTA-STS was done in three parts. First was checked if a domain had a DNS record for *mta-sts.domain.fi*. Second, if a record existed and pointed to an IP address, a HTTP request was made to download `mta-sts.txt` file from the path specified in the RFC: *https://mta-sts.domain.fi/.well-known/mta-sts.txt*. If the file was downloaded successfully, the validity of its format was checked.

Third part was to check if a domain had a valid TXT record for *_mta-sts.domain.fi*.

The implementation level of MTA-STS was even lower than the implementation level of DANE as can be seen in Table 13.

Table 13. MTA-STS statistics for *.fi* zone

Explanation	Count	Percentage	
Number of *.fi*-domains	**524728**		
DNS query timed out	747	0.14%	of *.fi* domains
no mta-sts A-record	446815	85.15%	of *.fi* domains
mta-sts A-record exists	77166	14.71%	of *.fi* domains
Couldn't resolve hostname	126	0.16%	of domains with mta-sts A-record
TCP-connection failed	3984	5.16%	of domains with mta-sts A-record
HTTP error code 400 or above	6654	8.62%	of domains with mta-sts A-record
Connection timeout	12729	16.50%	of domains with mta-sts A-record
SSL/TLS handshake failed	2192	2.84%	of domains with mta-sts A-record
No reply	1	0.00%	of domains with mta-sts A-record
Certificate not signed by trusted CA or name in certificate doesn't match with the service name	47677	61.78%	of domains with mta-sts A-record
other errors	1	0.00%	of domains with mta-sts A-record
`mta-sts.txt` file received	3802	4.93%	of domains with mta-sts A-record
Invalid `mta-sts.txt` file	3664	96.37%	of `mta-sts.txt` files received
Valid `mta-sts.txt` fle	**138**	3.63%	of `mta-sts.txt` files received
mode testing	33	23.91%	of valid `mta-sts.txt` files
mode enforce	105	76.09%	of valid `mta-sts.txt` files

The majority (over 85%) of fi-domains did not have an mta-sts.domain.fi record present in the DNS. An MTA-STS record was present in 14.71% of fi-domains (Figure 17).

Large number of certificate validation errors (47677) is explained by wildcard resource records in DNS: domains have set up a wildcard RR pointing to a server, but the certificate in that server doesn't have the same wildcard name in it's Common Name or Subject Alternate Name fields and that causes validation to fail.

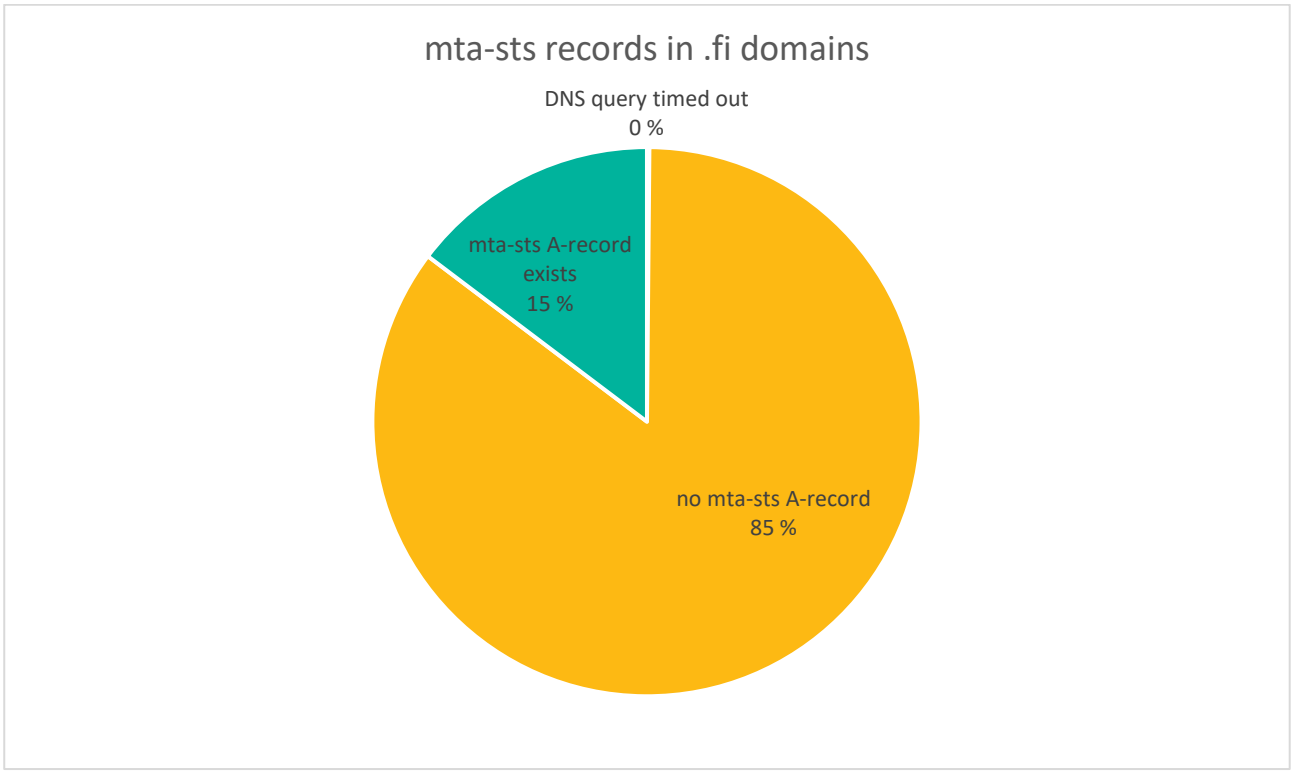

Figure 17. mta-sts records in *.fi* domains

Out of those, 3802 or 4.93%, provided a `mta-sts.txt` file over HTTPS (Figure 18).

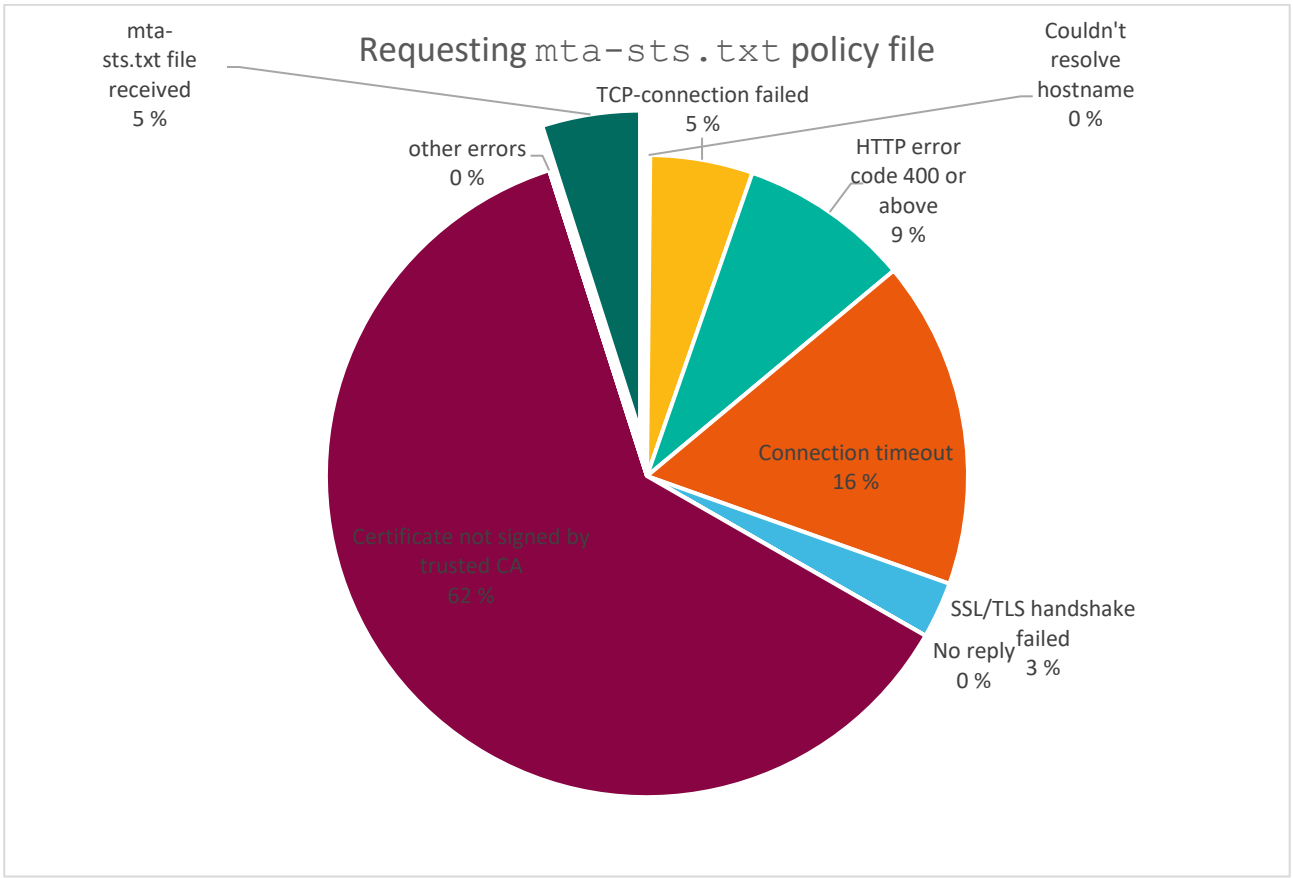

Figure 18. Requesting `mta-sts.txt` policy file

Most of the received files (over 96%) were incorrectly formatted (Figure 19). When looking at the results manually, it was clearly visible that the invalid files consisted mainly of different kind of place holders for reserved domains, HTTP-servers that replied with HTTP result code 200 to every request, and wild card redirects to organization's landing pages.

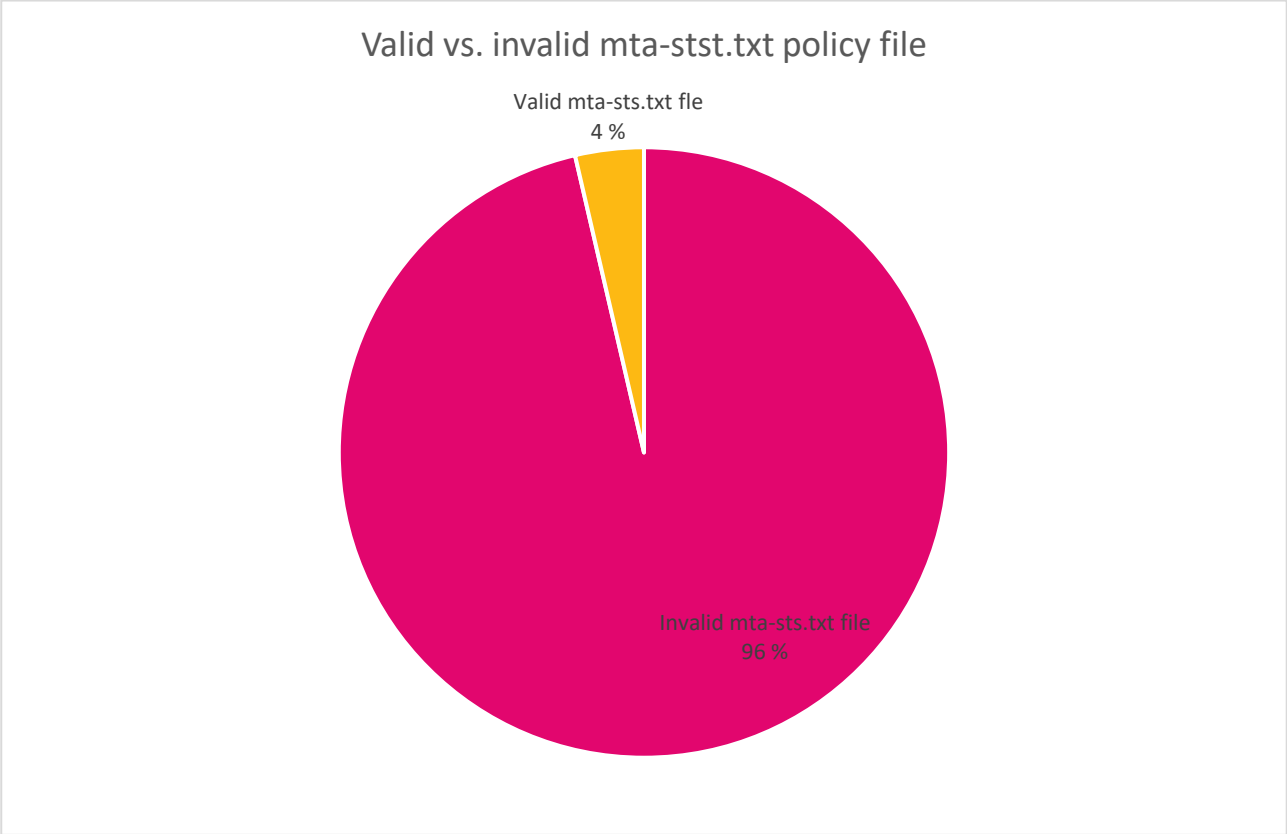

Figure 19. Valid vs. invalid `mta-sts.txt` policy file

Only 138 domains replied with a valid `mta-sts.txt` file and out of those, 33 were advertising testing mode and 105 enforcing mode (Figure 20).

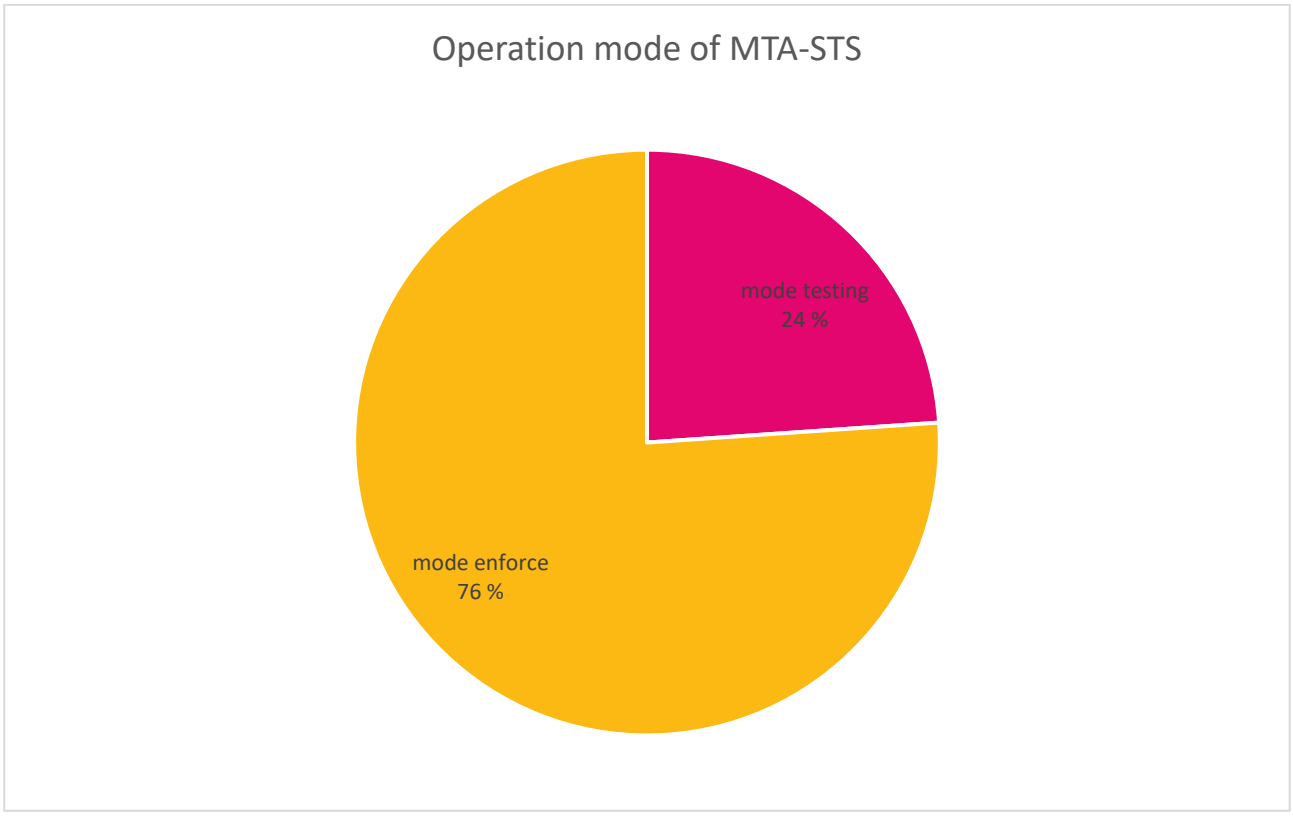

Figure 20. Operation modes of MTA-STS

In addition to a valid policy file, a valid _mta-sts TXT record is required in the DNS. Results can be seen in Table 14.

Table 14. _mta-sts TXT record results

Explanation	Count	Percentage	
_mta-sts TXT DNS query timeout	6	0.01%	of domains with existing mta-sts A-record
no _mta-sts TXT-record	60310	78.16%	of domains with existing mta-sts A-record
_mta-sts TXT-record exists	16850	21.84%	of domains with existing mta-sts A-record
valid _mta-sts TXT-record	**131**	0.78%	of domains with an _mta-sts record
invalid _mta-sts TXT-record	16719	99.22%	of domains with an _mta-sts record

Over 78% of domains with existing A-record for mta-sts do not have an _mta-sts record present in the DNS (Figure 21).

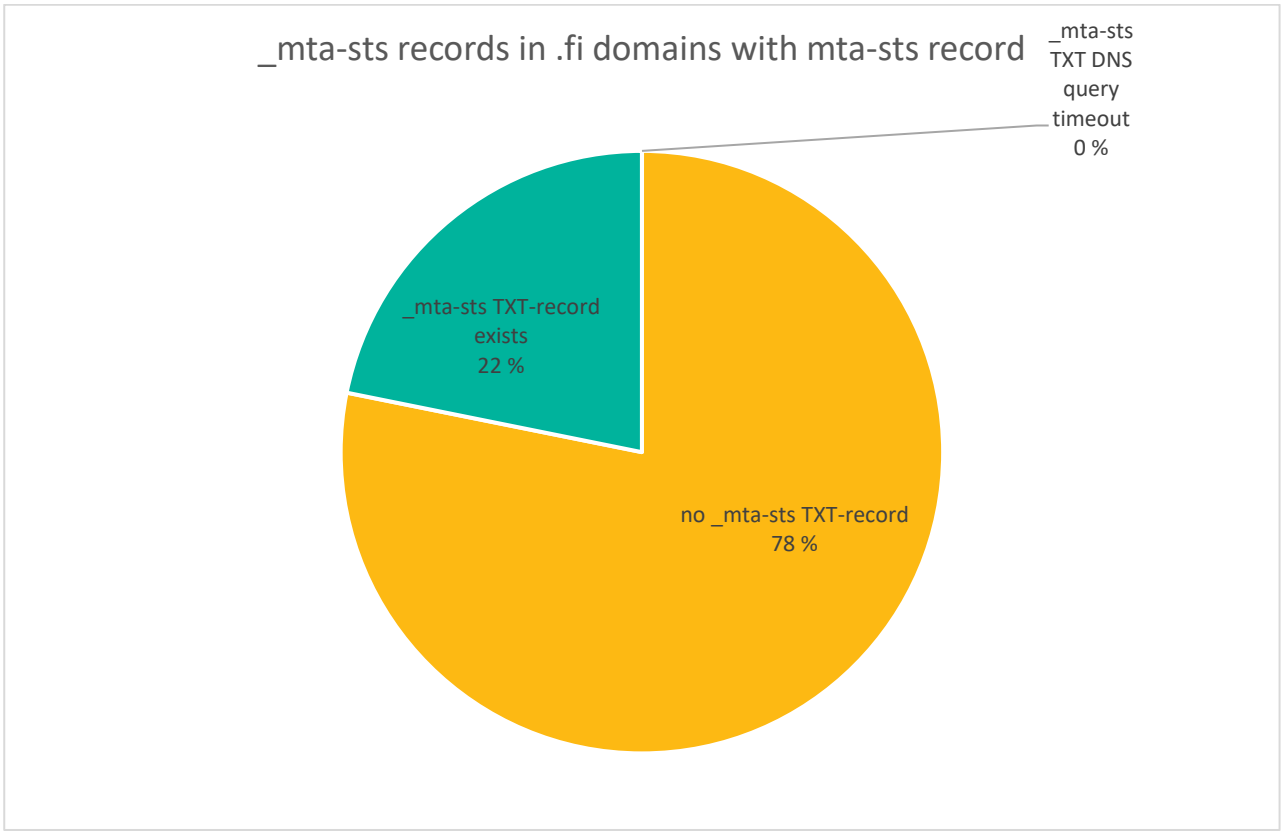

Figure 21. _mta-sts records in *.fi* domains with an mta-sts record

Out of all domains with existing _mta-sts record, only 131 have it in valid format (Figure 22).

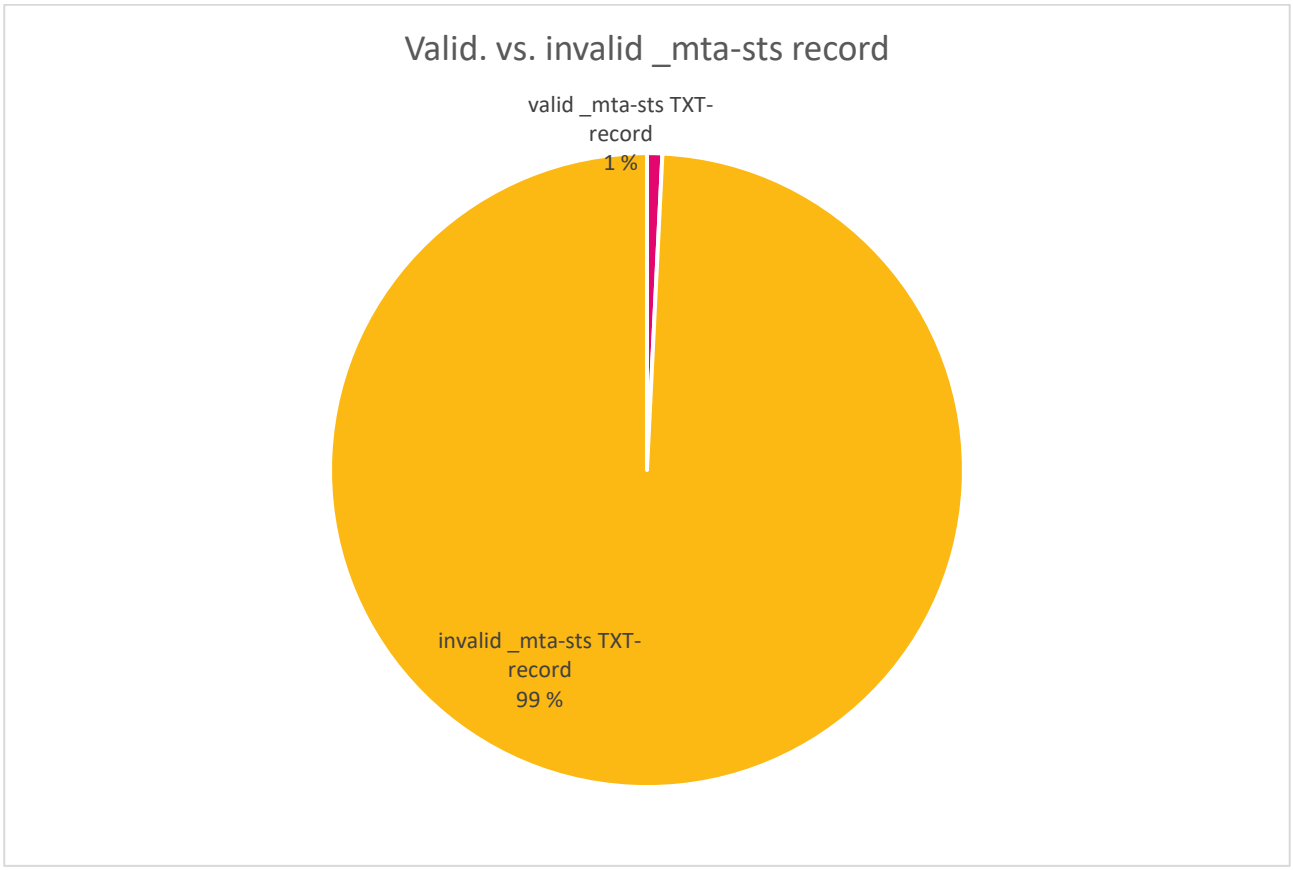

Figure 22. Valid vs. invalid _mta-sts record

MTA-STS needs a valid policy file provided from a webserver with a valid digital certificate signed by a trusted CA and a valid _mta-sts record with the correct version and ID. Combination of these results is shown in Table 15.

Table 15. fi domains with valid MTA-STS configuration

Explanation	**Count**	**Percentage**
Valid `mta-sts.txt` file	138	
Valid _mta-sts TXT record	131	
Only valid `mta-sts.txt` file	17	
Only valid _mta-sts TXT record	10	
Valid `mta-sts.txt` file AND _mta-sts TXT record	**121**	
out of all *.fi* domains		0.02%

As the results show, there are some domains with only a correct policy file or a correct _mta-sts record, but only 121 have them both properly configured. Key finding from these figures is that MTA-STS has not yet gained popularity among the *.fi* domains.

5.3 Combination of DANE and MTA-STS

DANE and MTA-STS can be used in parallel as they don't interfere with each other, and they provide additional protection compared to using only one of them.

By comparing fi-domains that have implemented DANE with domains that have implemented MTA-STS, the implementation level of both security controls in a single domain is very low: only 21 *.fi* domains have implemented them both as seen in Table 16.

Table 16. *.fi* domains with both DANE and MTA-STS implemented

Explanation	Count	Percentage
Domains with valid MTA-STS and DANE configuration	21	0.004% of all *.fi* domains

5.4 TLS-RPT

To reach the best possible protection and detection capability from attacks against SMTP encryption, a domain should have all three technologies configured.

From the domains that had DANE configured, only 10 also had TLS-RPT configured, as can be seen from Table 17.

Table 17. DANE and TLS-RPT configured

Explanation	Count	Percentage
Number of *.fi*-domains with DANE configured	4463	
Number of domains with TLS-RPT not configured	4453	99,78 % of *.fi*-domains with DANE
Number of domains with TLS-RPT configured	10	0,22 % of *.fi*-domains with DANE

From the domain that had MTA-STS configured, also 45 had TLS-RPT configured, as can be seen from Table 18.

Table 18. MTA-STS and TLS-RPT configured

Explanation	Count	Percentage
Number of *.fi*-domains with MTA-STS configured	121	
Number of domains with TLS-RPT not configured	75	61,98 % of *.fi*-domains with MTA-STS
DNS query timed out	1	0,83 % of *.fi*-domains with MTA-STS
Number of domains with TLS-RPT configured	45	37,19 % of *.fi*-domains with MTA-STS

From the domains that had both DANE and MTA-STS configured, 6 had also configured TLS-RPT as can be seen from Table 19.

Table 19. DANE, MTA-STS and TLS-RPT configured

Explanation	**Count**	**Percentage**
Number of *.fi*-domains with both DANE and MTA-STS configured	21	
Number of domains with TLS-RPT not configured	15	71,43 % of *.fi*-domains with MTA-STS and DANE
Number of domains with TLS-RPT configured	**6**	28,57 % of *.fi*-domains with MTA-STS and DANE

A key finding from this is that domains with MTA-STS configured are much more likely to also have TLS-RPT configured and that only 6 domains out of all 524728 *.fi*-domains in the scope of this study had all three security controls correctly configured.

6 Recommendations

6.1 Best practices

One of the goals within this study was to identify how to securely implement the researched techniques. Actual guidelines with step-by-step instructions about how to configure the protocols were not part of the assignment, but a top-level summarization is provided in the following subchapters.

To reach maximum protection from the threats discussed in chapter 3.3, all the following controls should be implemented.

6.1.1 DNSSEC

All domains receiving email should have DNSSEC enabled to protect from fake or counterfeit responses. This ensures the integrity of the DNS information that sending SMTP servers receive from other DNS servers, but does not directly improve confidentiality.

DNSSEC should also be enabled on all external domains hosting SMTP servers used to handle mail for the receiving domain.

6.1.2 DANE

TLSA records should be added to all SMTP server addresses that are referred to in MX records to protect against man-in-the-middle attacks, such as STRIPTLS and connection interception. Appropriate usage, selector and matching-type should be selected based on the digital certificate used on the SMTP server. TLSA records should also be present for all MX servers on external domains.

6.1.3 MTA-STS

MTA-STS should be configured on all domains receiving emails. In the policy file, the policy mode should be set to enforce, and all valid MX servers should be listed.

6.1.4 TLS-RPT

All domains that receive emails and that implement DANE and/or MTA-STS, should have TLS-RPT configured, so that the domain operators get reports about possible violations of policies. Using a service that interprets the JSON reports to a more human readable format makes monitoring easier.

6.1.5 Support for the protocols

Support for DANE and MTA-STS by SMTP servers and mail service providers is still limited, probably due to the low demand and the relatively young age of the protocols. If the domain operator is hosting their SMTP servers by themselves, they should choose an SMTP server that supports at least one of the protocols. If the domain operator is using an external service provider to host their mail, they should choose one that supports at least one of the protocols.

6.2 National guidelines and requirements

As the research shows, the implementation rate of the protocols is very low. To increase it, relevant authorities and interest groups should at first start to recommend using DNSSEC, DANE and MTA-STS and after suitable transition period start to demand using them at least from high-risk operators, like health care, financial sector and government institutions, bodies, and agencies.

7 Conclusions

7.1 Answers to the research questions

The research questions of this study were:

1. How to deploy DANE and MTA-STS securely?
2. What is their deployment rate in the *.fi* country code top-level domain (ccTLD)?

Both questions can be considered as answered. Based on studying the standards and documents analysis done on different guidelines a consensus was found on how to deploy the protocols securely. International guidelines support this conclusion, although the number of recommendations and guidelines published by recognized practitioners and official government authorities was limited.

The deployment rate of the protocols was found to be very low among the Finnish top-level domain *.fi*. Especially domains that had properly configured all controls was practically non-existent. There is a lot of work to be done to improve awareness of these modern protection mechanisms among domain operators. It is not likely that the usage of application layer encryption is going to increase among average users as it has its own difficulties, such as key management and application compatibility.

A good comparison point for DNSSEC and DANE implementation level was found to be the ccTLD of Netherlands (*.nl*). According to SIDN Labs, there are 6.2 million domains registered int the *.nl* ccTLD, out of which 3.6 million, or 58%, have DNSSEC enabled, compared to Finland, where only 3.45% of the *.fi* domains had DNSSEC enabled. DANE was implemented in 0.88 million, or 14% of the *.nl* domains, whereas in Finland only 4454 or 0.85% of *.fi* domains have DANE at least partially in use. One explanation to this is that the government authorities in Netherlands have taken active role in promoting the protocols and it is mandatory for government organizations to implement them.

Good comparison points for MTA-STS implementation rates were no found. Some guidelines recommend using it, but statistics of deployment rate were not available.

7.2 Reliability and ethicality

The starting point for assessing the implementation rate was an up-to-date zone file of *.fi* ccTLD received from the Finnish domain authority Traficom. Information was gathered from authoritative DNS servers and DNSSEC validation was in use and results were validated according to the standards. Based on these facts, the results can be considered reliable.

As this was only a one-time study, it will only show the situation at one point-in-time and does not provide information about any changes in the implementation numbers. But considering the nature of the protocols covered in the study, rapid changes can happen when large service providers implement the protocols, but constant fluctuation is not likely.

All information used in the research except the *.fi* zone file is publicly available online. Individual domain names in the zone file are not confidential as such and can be gathered for example by using brute force enumeration. Traficom even has an API that anyone can query to find domain names and information related to them based on different parameters. But the whole zone file is not published as it might be used as a target list by a threat actor.

Because the majority of the *.fi* domains do not use the protocols discussed in this study and that is also the presumption, results of this study can't be considered increasing risk to Finnish domain owners.

No abuse contacts were received during the first month after the information gathering was done.

7.3 Discussion

Considering that many governments are monitoring both domestic and international communications to gather intelligence information to improve their national security, to gain commercial benefit, or to monitor their own citizens to keep them under control, the implementation rate in Finland is surprisingly low.

This study concentrated on the whole *.fi* ccTLD and not on different interest groups, so no comparison on the implementation rate between different groups of *.fi* domains was done. A topic for further research could be to select interest groups that could be considered as high-risk targets and do an action study research to see if the implementation rate can be increased with guidelines or requirements from the authorities.

Keeping up with the ever-increasing risks and threats, cyber security needs continuous monitoring and improvement. The study has shown that there is room for improvement from both the authorities and individual domain owners to raise awareness, develop regulation and implement security controls to protect confidentiality, integrity, and availability of email communications.

References

Australian Cyber Security Centre. 2022. Guidelines for Email. Accessed on 25.4.2022. Retrieved from https://www.cyber.gov.au/acsc/view-all-content/advice/guidelines-email

Bursztein E. 2015. Understanding how tls downgrade attacks prevent email encryption. Accessed on 9.5.2022. Retrieved from https://elie.net/blog/understanding-how-tls-downgrade-attacks-prevent-email-encryption/

Canadian Centre for Cyber Security. 2021. Implementation guidance: email domain protection (ITSP.40.065 v1.1). Accessed on 26.4.2022. Retrieved from https://www.cyber.gc.ca/en/guidance/implementation-guidance-email-domain-protection

Cloudflare. 2022. What is BGP hijacking? Accessed on 25.4.2022. Retrieved from https://www.cloudflare.com/learning/security/glossary/bgp-hijacking/

DNSSEC-tools. 2022. DNSSEC and DANE Deployment Statistics. Accessed on 27.4.2022. Retrieved from https://stats.dnssec-tools.org/

European Commission Joint Research Centre (JRC). 2015. A Security Analysis of email communications.

Forum Standaardisatie. 2022. Verplichte standaarden. Accessed on 10.5.2022. Retrieved from https://www.forumstandaardisatie.nl/open-standaarden/verplicht

Google. 2022. Email encryption in transit. Accessed on 27.4.2022. Retrieved from https://transparencyreport.google.com/safer-email/overview?hl=en

Hoffman P. 2002. SMTP Service Extension for Secure SMTP over Transport Layer Security. Accessed on 1.4.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc3207

Internet Engineering Task Force (IETF). 2015a. SMTP Security via Opportunistic DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS). Accessed on 18.4.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc7672

Internet Engineering Task Force (IETF). 2015b. A "Null MX" No Service Resource Record for Domains That Accept No Mail. Accessed on 24.4.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc7505

Internet Engineering Task Force (IETF). 2018a. SMTP MTA Strict Transport Security (MTA-STS). Accessed on 18.4.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc8461

Internet Engineering Task Force (IETF). 2018b. SMTP TLS Reporting. Accessed on 3.5.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc8460

Klensin J. 2008. Simple Mail Transfer Protocol. Accessed on 20.4.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc5321

Kontinen Ville. 2020. Master's thesis: Preventing email forgery in Finland - Research on the current SPF and DMARC implementations. JAMK University of Applied Sciences

Microsoft. 2022. Introducing MTA-STS for Exchange Online. Accessed on 27.4.2022. Retrieved from https://techcommunity.microsoft.com/t5/exchange-team-blog/introducing-mta-sts-for-exchange-online/ba-p/3106386

Mockapetris P. 1987. Domain Names – Implementation and Specification. Accessed on 20.4.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc1035

National Cyber Security Centre. 2019. Email security and anti-spoofing. Accessed on 26.4.2022. Retrieved from https://www.ncsc.gov.uk/collection/email-security-and-anti-spoofing/using-mta-sts-to-protect-the-privacy-of-your-emails

Network Working Group (NWG). 2005. DNS Security Introduction and Requirements. Accessed on 20.4.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc4033

Network Working Group (NWG). 2007. OpenPGP Message Format. Accessed on 24.4.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc4880

Network Working Group (NWG). 2009. Cryptographic Message Syntax (CMS). Accessed on 24.4.2022. Retrieved from https://datatracker.ietf.org/doc/html/rfc5652

Postel, J. 1982. Simple Mail Transfer Protocol. Accessed on 24.3.2022. Retrieved from https://tools.ietf.org/html/rfc821

SIDN. 2019. Number of DANE-enabled mail domains growing exponentially. Accessed on 16.5.2022. Retrieved from https://www.sidn.nl/en/news-and-blogs/number-of-dane-enabled-mail-domains-growing-exponentially

SIDN Labs. 2022. *.nl* statistics. Accessed on 27.4.2022. Retrieved from https://stats.sidnlabs.nl/en/mail.html

Statista. 2022. Number of e-mail users worldwide from 2017 to 2025. Accessed on 24.4.2022. Retrieved from https://www.statista.com/statistics/255080/number-of-e-mail-users-worldwide/

The Record. 2022. Russia reroutes internet in occupied Ukrainian territory through Russian telcos. Accessed on 3.5.2022. Retrieved from https://therecord.media/ukraine-internet-blackout-kherson-skynet-russia/

Turunen Matti. Master's thesis: State of email security implementations in Finnish municipalities and joint municipal authorities in 2021. JAMK University of Applied Sciences

UK Government (gov.uk). Set up government email services securely. 2021. Accessed on 25.4.2022. Retrieved from https://www.gov.uk/guidance/set-up-government-email-services-securely

Appendices

Appendix 1. Sample plain text SMTP transmission

```
1: 220 mail.domainA.fi ESMTP OpenSMTPD

2: HELO smtp.domainB.fi

3: 250 mail.domainA.fi Hello smtp.domainB.fi [192.168.0.254],
pleased to meet you

4: MAIL FROM:<sender@domainB.fi>

5: 250 2.0.0 Ok

6: RCPT TO:<receiver@domainA.fi>

7: 250 2.1.5 Destination address valid: Recipient ok

8: DATA

9: 354 Enter mail, end with "." on a line by itself

10: MIME-Version: 1.0

    Content-Type: text/plain; charset=utf-8

    Content-Transfer-Encoding: quoted-printable

    Subject: Test message

    From: Sender Name <sender@domainB.fi>

    To: receiver@domainA.fi

    X-Priority: 5

    Content of the message

    .

11: 250 2.0.0 6b0b314e Message accepted for delivery

12: QUIT
```

```
13: 221 2.0.0 Bye
```

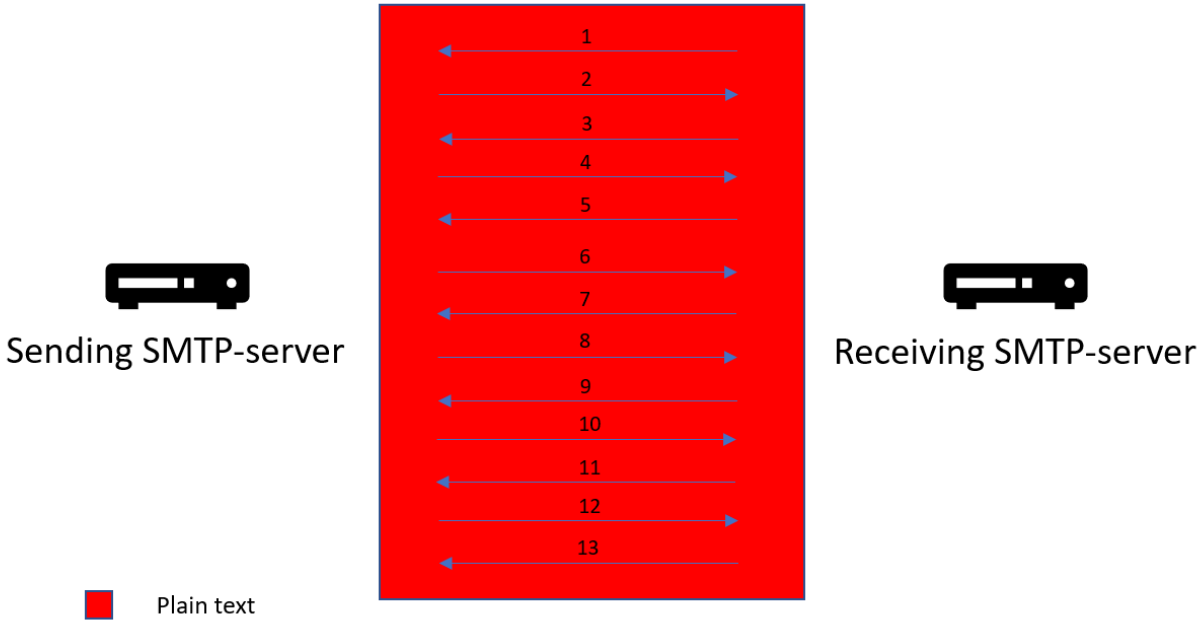

Figure 23. Sample plain text SMTP transmission

Appendix 2. Sample SMTP transmission with STARTTLS and encryption

```
1: 220 mail.domain.fi ESMTP OpenSMTPD
2: EHLO EUR05-VI1-obe.outbound.protection.outlook.com
3: 250-mail.domain.fi Hello EUR05-VI1-
obe.outbound.protection.outlook.com [40.107.21.92], pleased to
meet you
   250-8BITMIME
   250-ENHANCEDSTATUSCODES
   250-SIZE 36700160
   250-DSN
   250-STARTTLS
   250 HELP
4: STARTTLS
5: 220 2.0.0 Ready to start TLS
```

- - - TLS Handshake Starts - - -

6: *Client Hello*

7: *Server Hello*

Certificate, Server Key Exchange, Server Hello Done

8: *Client Key Exchange, Change Cipher Spec, Encrypted Handshake Message*

9: *Change Cipher Spec, Encrypted Handshake Message*

- - - TLS Handshake Finished - - -

10: *Encrypted message data*

11: *Acknowledgement of receiving the message*

12: *Closing the connection*

…

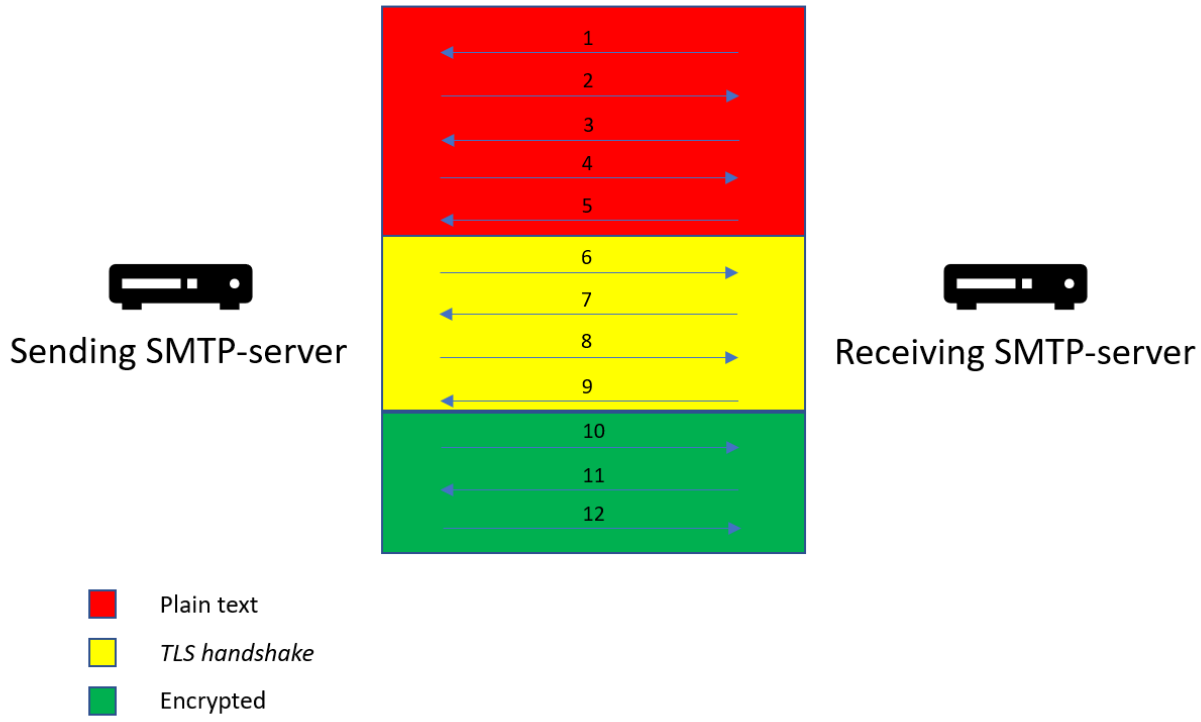

Figure 24. Sample SMTP transmission with STARTTLS and encryption