OARC 37: Internet Survey of DANE/TLSA DNS Records

Application and Use for Mail Exchanger (MX)

Erwin Hoffmann

[feh@fehcom.de]

Frankfurt University of Applied Sciences

February, 17th 2022

Motivation and Research Question

Background:

Since about 1998 I maintain, install, refactor, and enhance Daniel J. Bernstein's now public domain software. In particular

- **qmail** → **s/qmail**
- **djbdns** → **djbdnscurve6**

to include IPv6 capabilities, TLS encryption, and other missing features. Partially as reseach and teaching playground at my University, and partially to supply the still-existing DJB community with today's required solutions.

TLS and DANE for Mail Exchanger:

Having implemented TLS for **qmail** and establishing the fork **s/qmail**, in version 4.1 (apart from native SPF support) I've included TLSA lookups for its mail client – **qmail-remote** – based upon a new DNS stub resolver library providing some additional command-line tools:

- **dnsmxip** – lookup of MX record and the adjacent IPv6/IPv4 addresses.
- **dnstlsa** – lookup of TLSA records for the given MX (while synthesizing the TLSA FQDN): `mx.example.com` → `_25._tcp.mx.example.com`

In addition, the DNSCurve enabled authoritative DNS server **tinydns** supports TLSA/DANE (and DKIM) records out-of-box.

Research Questions

- To which extend are TLSA records deployed in the DNS in particular for MX services (on TCP port 25)?
- Since DANE allows four *Usage* modes, what is the current situation here?
- Which other constraints (*Selector*, *Matching Types*) are actually realized within TLSA records?
- How do people handle key-rollover of X.509 certs and availability?
- What are the usual operational conditions for DANE records to support TLS services for MX servers?

↪ By means of an extensive DNS Internet survey those questions were tackled.

Out-of-scope:

- The combination of TLSA records and DNSSec.
- Whether the TLSA record matches the presented X.509 certificate in the TLS handshake indeed.

Based upon zone data from [`https://zonefiles.io/`] within the Bachelor thesis of *Jihad El Hayek* we queried in Q4/2021 the following gTLD and ccTLD domains:

- gTLD: `INFO`, `ORG`, `NET`
- Europe: `AT`, `BE`, `CH`, `CZ`, `DE`, `ES`, `EU`, `FR`, `IT`, `PL`, `UK`, `RU`, `SE`
- America, Asia & Pacific: `CA`, `BR`, `AU`, `CN`, `JP`, `NZ`

↪ Note: The individual zone data were *incomplete* (given today's knowledge)!

Setup

Our setup was a two-tier:

A) MX qualification: For each TLD all included domains were subject of a MX lookup:

- Domains without MX or irregular MX entries like '`localhost`' were omitted.
- Domains pointing with their MX services to '`google.com`', '`hotmail.com`', or '`outlook.com`' were not considered as well.
- For each domain, a quadratic DNS lookup was performed, mimicking the usual query by an (E)SMTP client.

↪ The total set of FQDNs was about 66 mio. domains which were included in this step.

B) DANE qualification: Received valid MX responses became subject of a TLSA query:

- DNSSec was again not required here.
- **dnscache** did a iterative query on the authoritative NS of the respective zone.
- UDP and TCP DNS responses were considered.
- The UDP buffer for DNS messages was given by the (IPv6) `MTUSIZE-52`, and thus `1228` bytes. EDNS(0) responses were accepted as well.

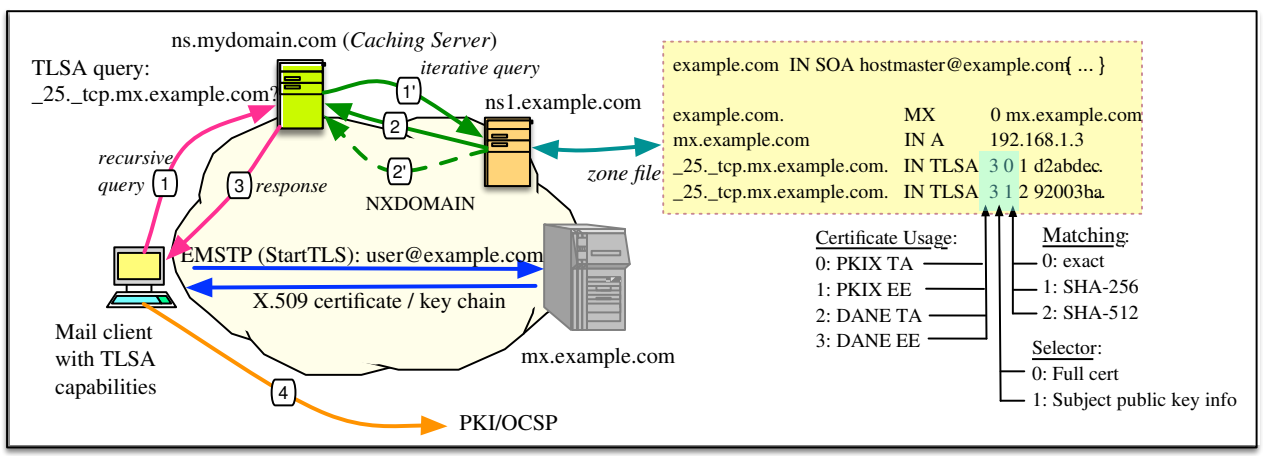

Figure: Typical setting of the DANE/TLSA query/response chain.

In total > 300 mio. DNS queries were performed. The set of received DNS responses (about 1 Gbyte) is available on GitHub for further analysis.

First Step: Received Bulk MX Data

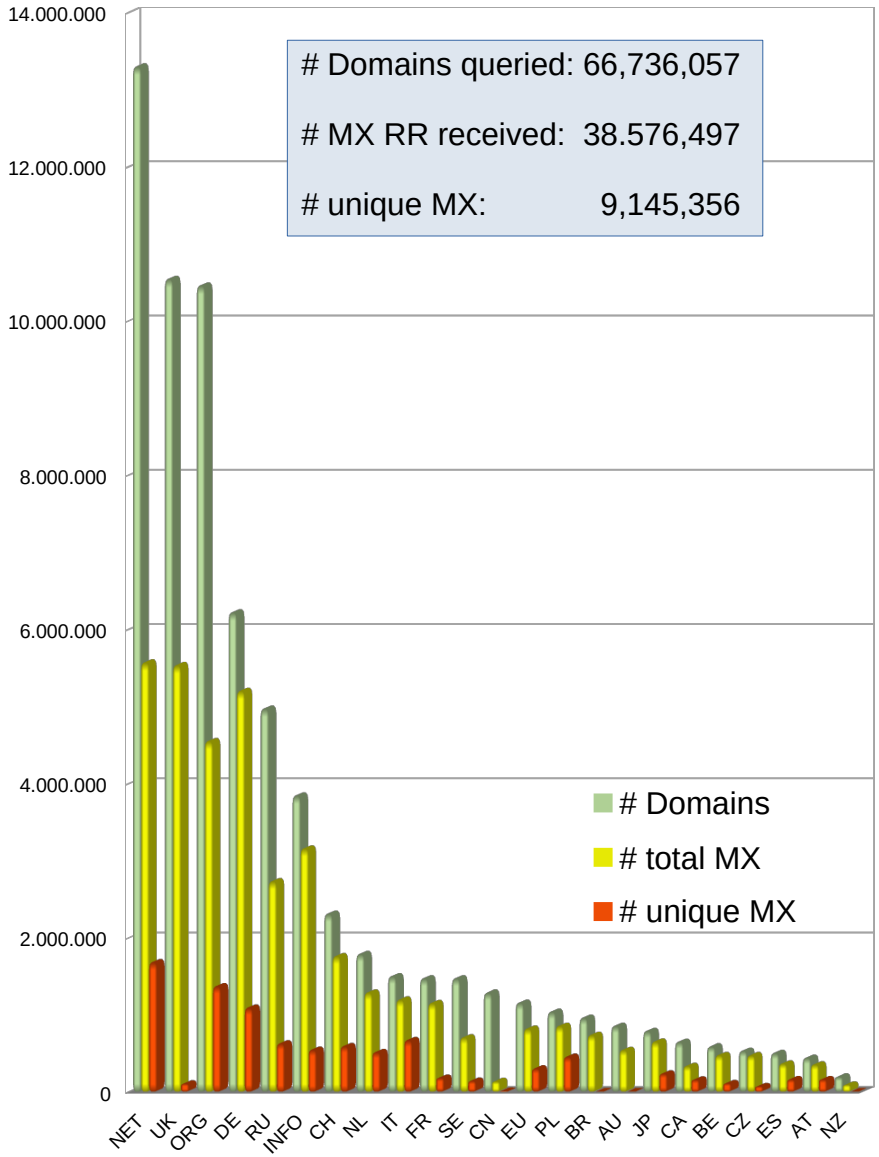

Figure: Number of queried domains per TLD, number of received MX records, and unique MX.

Query coverage compared to [8]:

- Totals: 66.7 mio out of 363 mio (18.5%)
- DE zone: 6 mio out of 12 mio.
- EU zone: 1.1 mio out of 3.6 mio.

Samples of MX responses:

```
2021-12-15 21:13:01 rr 55e9a044
86400 mx speedspharmacy.co.uk.  1
speedspharmacy-co-uk.mail.protection.outlook.com.

2021-12-15 21:14:55 rr 42608ea2 3600 mx
publimetrics.net.  30 mx.publimetrics.net.
2021-12-15 21:14:55 rr 42608ea2 3600 mx
publimetrics.net.  1 aspmx.l.google.com.
2021-12-15 21:14:55 rr 42608ea2 3600 mx
publimetrics.net.  10 alt4.aspmx.l.google.com.

2021-12-15 21:25:33 rr 5bc3f008 3600 mx
publisher.net.  0 localhost.

2021-12-15 21:26:03 rr c0ae440a 300 mx
stayzone.org.  0 .
```

Definition

Unique MX:

MX with distinguished FQDN.

↪ **MX delegation**! → MXaaS

Second Step: TLSA Responses

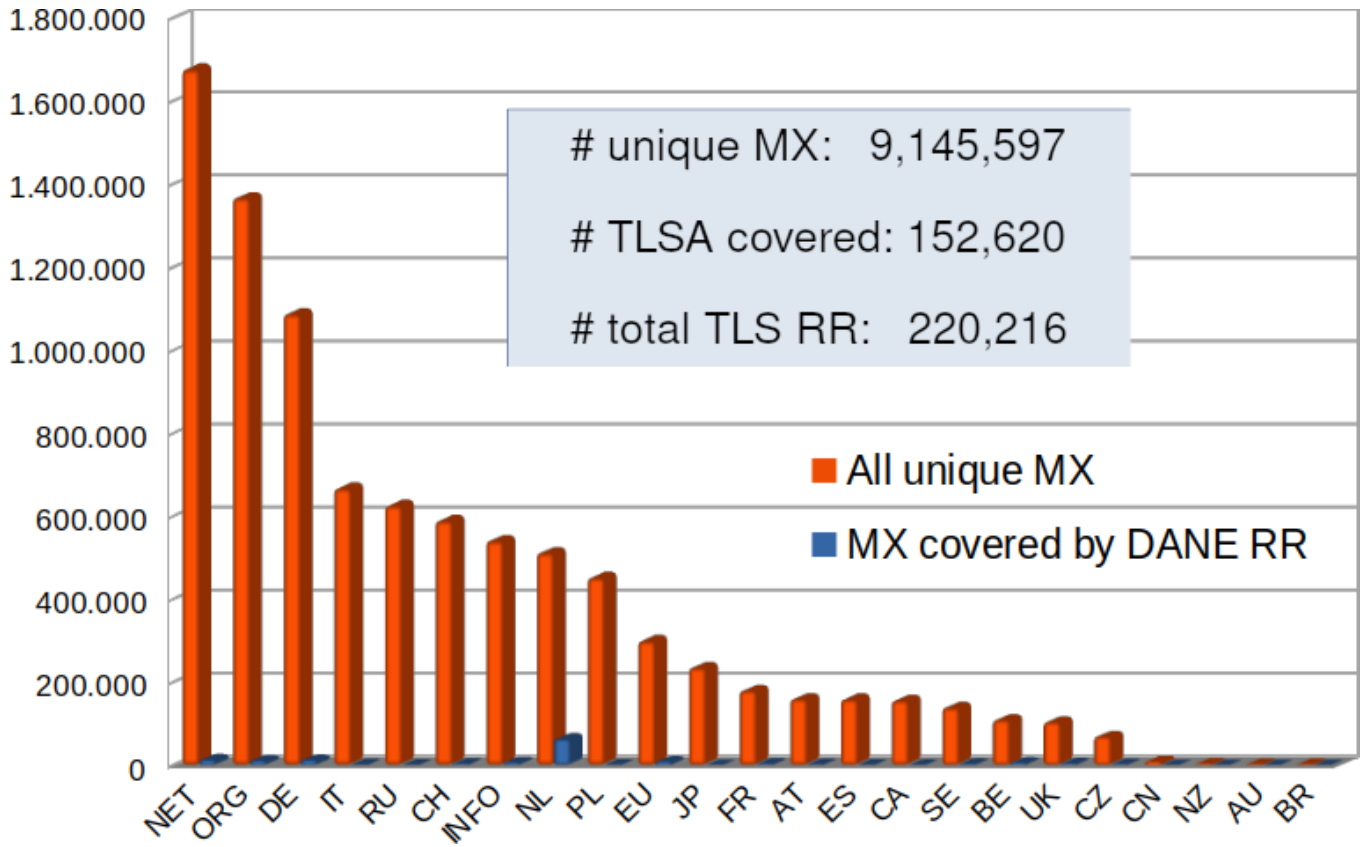

Figure: Number of uniquely received MX per TLD and TLSA coverage.

↪ Remember: The receiving MX might not be part of the queried domain name!

TLSA Responses - Closer Look

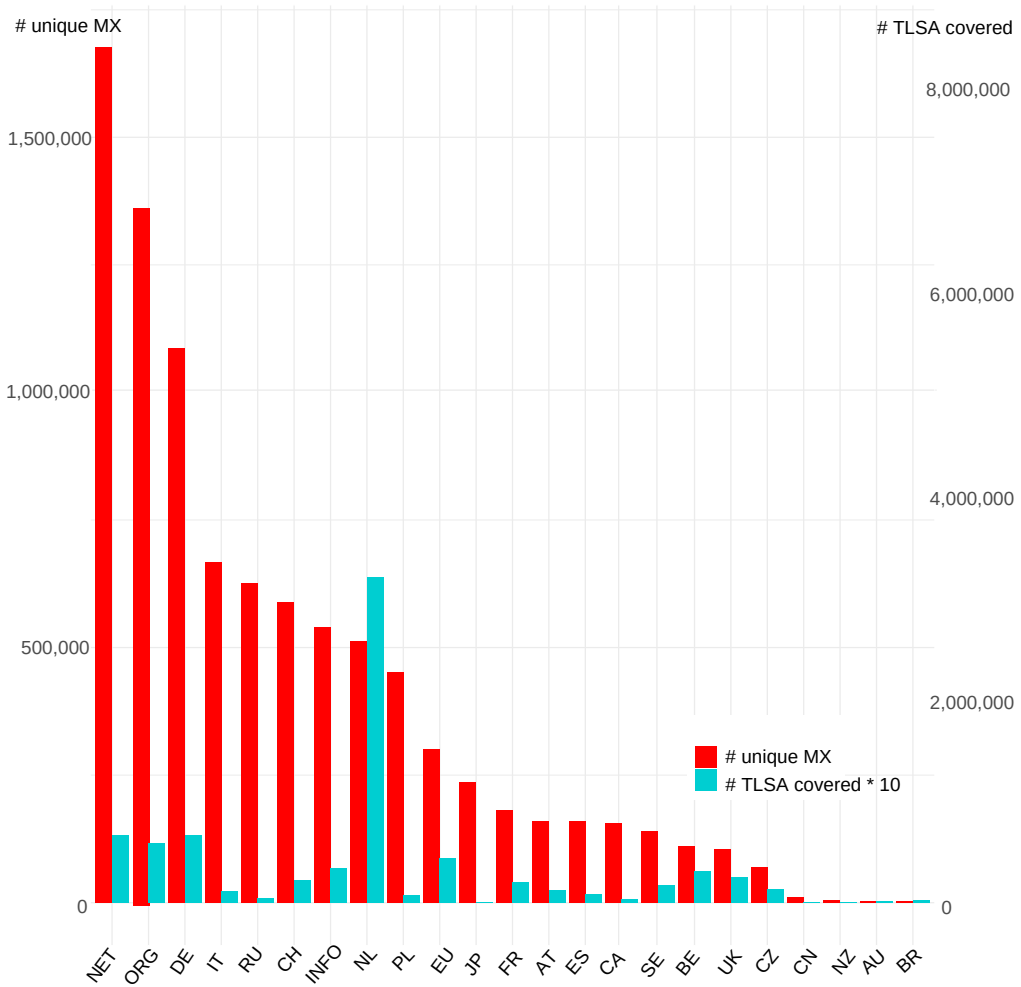

Figure: Number of uniquely received MX per TLD and TLSA coverage $\times$ 10.

$\hookrightarrow$
- The TLSA average coverage per unique MX is 1.67%,
- the average coverage per domain 0.38% (only)!

Detail Analysis of TLSA Results – Responses

In our further analysis we investigated which information is present in the response. We get:

- The *Usage*: (0) `PKIX-TA`, (1) `PKIX-EE`, (2) `DANE-TA`, (3) `DANE-EE`.
- The *Selector*: (0) Fingerprint of entire X.509 cert, (1) Fingerprint of *Subject Public Key Identifier* (**SPKI**).
- The *Matching Type*: (0) Entire X.509 cert, (1) SHA-256 digest, (2) SHA-512 digest.
- The *fingerprint* of (or the entire) X.509 certificate.
- The *number* of TLSA responses provided.

Sample of TLSA responses (**dnstlsa**) for the AU domain:

```
box.ninkeri.com.au=
Usage:  [3], Selector:  [1], Type:  [1] e0e8272da8b3ecb7f820aaeb85be00e4e4c4fc552b864f39ef1e7ea2fd1c9a

box.ocollins.me=
Usage:  [3], Selector:  [1], Type:  [1] 40d751567dd5e1e5d6bcf6c5ddae3ff5ccb1c78cc7a4b1d9f7dfc1b5b8792f

mail.naturaltherapiesandbeauty.com.au=
Usage:  [3], Selector:  [0], Type:  [1] e6b216a0166da9ab95c0b509585413c53f238300068cdfd2a43f66d8d0fb11

mail.protonmail.ch=
Usage:  [3], Selector:  [1], Type:  [1] 76bb66711da416433ca890a5b2e5a0533c6006478f7d10a4469a947acc8399
Usage:  [3], Selector:  [1], Type:  [1] 6111a5698d23c89e09c36ff833c1487edc1b0c841f87c49dae8f7a09e11e97
```

TLSA Responses – Usage

The *Usage* simply says,

a) whether the X.509 certificate given its MX it covers, is a PKIX derived cert (signed by a CA) (PKIX-TA, PKIX-EE), or

b) the X.509 certificate is self-signed (DANE-TA, DANE-EE).

Additionally, it can be told, whether the entire certificate chain is present during the TLS handshake (-TA), or just the 'endpoint' X.509 cert (-EE). Analyzing the given answers in TLSA records, the *Usage* shows a clear winner:

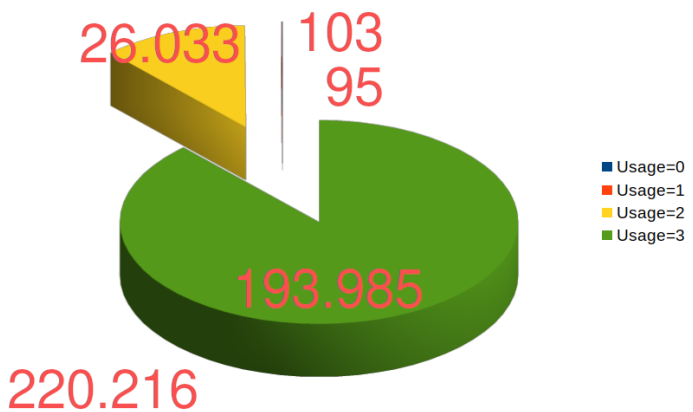

Figure: Breakdown of *Usage* for the received TLSA records; Usage (3): 88.09%, Usage (2): 11.82%.

↪ Operators of *Mail Exchangers* only see a need to deploy TLSA records in case of self-signed X.509 certificates – without the need for intermediate certs.

TLSA Responses – Selector

Both, *Selector* and *Matching Type* are operational parameters only. Here, *Selector* tells which part of the X.509 is covered, where only two solutions are possible:

- The standard fingerprint of the X.509 certificate, being subject of change after each renewal.
- The fingerprint of its *Subject Public Key Identifier* **SPKI** only; staying constant over cert exchanges.

↪ Thus, from a operational point of view the latter is immune against X.509 certificate roll-over: The DANE information may stay constant and does not conflict with certificate renewal. It is no surprise, that this is a preferred option for DANE:

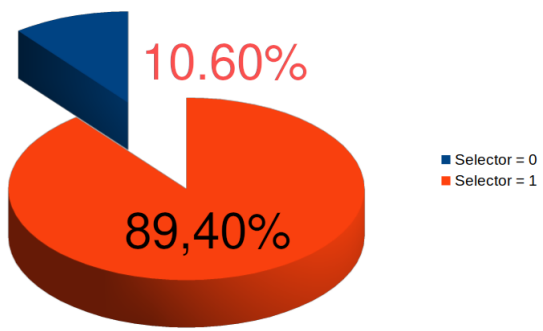

Figure: Breakdown of the *Selector* in TLSA records for MX.

TLSA Responses – Matching Type

We recall, that the *Matching Type* tells, how the information present in the TLSA Record shall be used for comparison:

- Type 0: The full certificate ($\sim$ 1 kbyte).
- Type 1: The SHA-256 hashsum (32 byte).
- Type 2: The SHA-512 hashsum (64 byte).

$\hookrightarrow$ There is not doubt, that MX and DNS operators tend to prefer SHA-256 hashsums; in particular considering the additional data present in DNSSec responses.

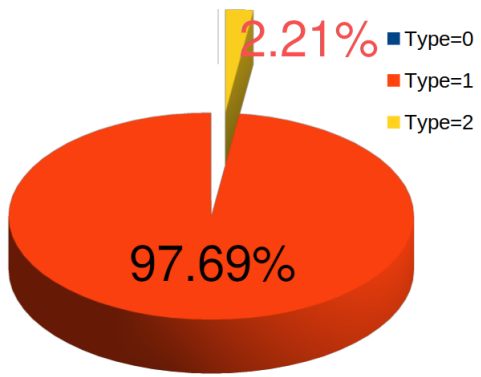

Figure: Breakdown of the *Matching Type* in TLSA records for MX.

TLSA Responses – How many Certs

Our final questions targets how many certs are provided per MX (in average). There may be two reasons to provide multiple responses here:

- The very same X.509 certificate is covered by different fingerprint and/or hashsums.
- One takes care of a certificate renewal and define multiple entries in here, covering retired, current, and potentially future valid certificates.

The answer is given here:

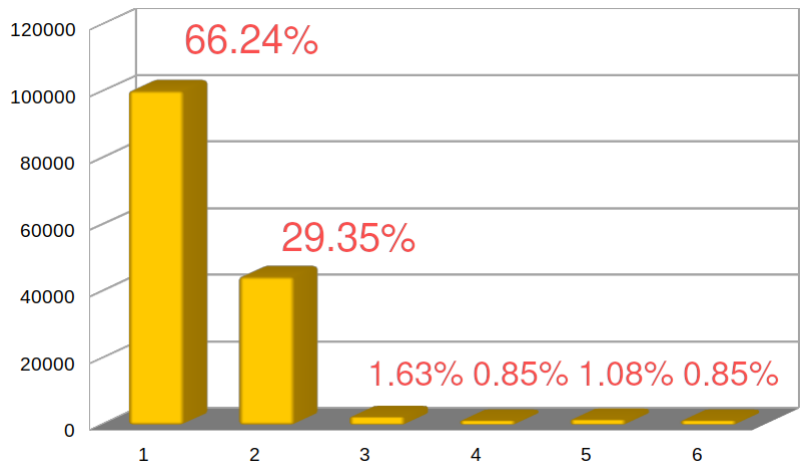

Figure: Breakdown of the *Matching Type* in TLSA records for MX.

↪ Thus, about 1/3 domains provide more than one TLSA record pointing to the same or different X.509 certificates. The average value here is 1.44.

Final Observations and Conclusion

- Given MX records for domains we see a huge 'MX delegation'.
- In particular the MX for Google, Microsoft (`hotmail.com`, `outlook.com`) but also smaller providers like `mailbox.org`, and `protonmail.ch` are often used to receive and queue (E)SMTP mails (on the same MX) for a lot of providers (and receiving them in clear text, even if TLS is used).
- TLSA coverage thus is strongly correlated with the fact that any of the major MX players provide TLSA records in the DNS or not.
- For instance, all (few) China (`CN`)e TLSA records are due to MX delegation.
- The overall coverage of DANE/TLSA records is still marginal; except for the Dutch domain (`NL`).
- The *Usage* policy of the TLSA records does not enforce the mail clients to use PKIX verification however permitting *self-signed* (DANE) X.509 certificates, providing their *Subject Public Key Identifier* SPKI by a SHA-256 hashsum.
- Differences among the TLDs regarding the TLSA operations are negligible; thus follow the same (practical) spirit.

Our results complement the periodic TLSA survey from *Viktor Dukhovni*, author of RFC 7671.

Final Observations and Conclusion

- Given MX records for domains we see a huge 'MX delegation'.
- In particular the MX for Google, Microsoft (`hotmail.com`, `outlook.com`) but also smaller providers like `mailbox.org`, and `protonmail.ch` are often used to receive and queue (E)SMTP mails (on the same MX) for a lot of providers (and receiving them in clear text, even if TLS is used).
- TLSA coverage thus is strongly correlated with the fact that any of the major MX players provide TLSA records in the DNS or not.
- For instance, all (few) China (`CN`)e TLSA records are due to MX delegation.
- The overall coverage of DANE/TLSA records is still marginal; except for the Dutch domain (`NL`).
- The *Usage* policy of the TLSA records does not enforce the mail clients to use PKIX verification however permitting *self-signed* (DANE) X.509 certificates, providing their *Subject Public Key Identifier* SPKI by a SHA-256 hashsum.
- Differences among the TLDs regarding the TLSA operations are negligible; thus follow the same (practical) spirit.

Our results complement the periodic TLSA survey from *Viktor Dukhovni*, author of RFC 7671.

Thank you for your attention!

Final Summary of your Analysis

Table: Total evaluated domains and received MX and TLSA RR.

	Number of ...			MX with TLSA		Usage				Selector		Matching Type		
Domain	Domains	covered MX	unique MX	covered	RRs	0	1	2	3	0	1	0	1	2
INFO	3,836,691	3,151,408	540,424	6,761	10,323	15	7	1,877	8,424	1,455	8,870	8	9,946	377
NET	13,292,521	5,562,618	1,675,649	13,195	19,253	19	21	3,421	15,776	2,869	16,377	5	18,536	709
ORG	10,450,626	4,542,621	1,366,267	11,539	17,420	29	16	3,099	14,244	2,996	14,402	16	16,717	671
BR	955,264	733,387	2,458	464	783	0	0	16	748	45	738	0	768	15
CA	643,077	330,883	15,6184	653	1,215	0	0	168	1,047	80	1,135	0	1,214	1
AT	436,132	351,897	159,878	2,366	4,058	1	1	475	3,581	375	3,683	0	3,974	84
BE	583,142	471,817	109,734	6,146	8,764	0	0	1,347	7,471	812	7,952	0	8,538	226
CH	2,304,768	1,753,310	589,175	4,270	6,675	8	7	962	5,698	1,217	5,458	0	6,456	219
CZ	526,665	469,462	69,356	2,631	3,806	0	0	407	3,399	300	3,506	0	3,705	101
DE	6,213,317	5,195,049	1,086,378	13,114	20,123	7	21	3,349	16,646	5,249	14,874	0	19,264	856
ES	498,069	363,879	159,191	1,606	2,860	0	0	265	2,595	203	2,657	0	2,830	30
EU	1,149,377	814,803	299,683	8,693	12,763	2	5	2,059	10,697	1,893	10,302	1	12,302	460
FR	1,466,033	1,145,294	180,188	3,946	5,826	0	0	780	5,046	568	5,258	0	5,687	139
IT	1,493,905	1,191,211	667,064	2,258	3,785	0	0	387	3,388	356	3,429	3	3,722	63
NL	1,782,553	1,282,416	511,681	63,715	83,918	3	0	4,496	79,419	3,091	80,827	0	83,195	723
PL	1,032,558	85,0241	452,154	1,493	2,555	0	1	178	2,376	239	2,316	0	2,495	60
RU	4,962,992	2,726,402	625,221	899	1,455	0	0	335	1,120	175	1,280	0	1,423	32
SE	1,465,433	702,868	138,956	3,415	5,758	0	0	753	5,005	444	5,314	0	5,657	101
UK	10,536,401	5,527,898	104,687	4,853	7,840	19	16	1,302	6,483	1,005	6,815	12	7,644	170
AU	849,744	53,921	2,808	294	517	0	0	109	408	98	419	0	515	2
CN	1,280,957	134,696	9,342	87	133	0	0	19	114	6	127	0	132	1
JP	78,3319	646,255	235,005	96	209	0	0	61	148	11	198	0	209	0
NZ	19,2513	8,8791	3,873	126	235	0	0	29	206	28	207	0	234	1

Sources & References


Hoffmann, E.: *djbdnscurve6*, https://www.fehcom.de/ipnet/djbdnscurve6.html


Hoffmann, E.: *s/qmail*, https://www.fehcom.de/sqmail/sqmail.html


Hoffmann, E.: *DNS TLSA Survey*, https://github.com/ErwinHo/DNS_TLSA_Survey/releases/tag/v1.0


Barnes, R.: *RFC 6394: Use Cases and Requirements for DNS-Based Authentication of Named Entities (DANE)*, https://datatracker.ietf.org/doc/html/rfc6394


Dukhovni, V. and Hardaker, W.: *RFC 7671: The DNS-Based Authentication of Named Entities (DANE) Protocol: Updates and Operations Guidance*, https://datatracker.ietf.org/doc/html/rfc7671


Zhu, l., Wessels, D., Mankin, A. and Heidmann, J.: *Measuring DANE TLSA Deployment*, https://link.springer.com/chapter/10.1007%2F978-3-319-17172-2_15


Dukhovni, V.: *Update on stats 2021-11*, https://www.mail-archive.com/dane-users@sys4.de/msg00473.html


European Commission: *Study on Domain Name System (DNS) Abuse*, https://op.europa.eu/en/publication-detail/-/publication/7d16c267-7f1f-11ec-8c40-01aa75ed71a1/language-en/