

DNS-Based Authentication of Named
Entities (DANE)
Internet-Draft
Intended status: Standards Track
Expires: December 29, 2012

T. Finch
University of Cambridge
June 27, 2012

Secure SMTP with TLS, DNSSEC and TLSA records.
draft-fanf-dane-smtp-04

Abstract

SMTP has a STARTTLS extension, but (especially in the case of inter-domain mail transfer) it only provides very limited security because it does not specify how to authenticate the server's certificate. This memo specifies how TLSA records in the DNS can be used for proper SMTP server authentication.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on December 29, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as

described in the Simplified BSD License.

Table of Contents

1. Introduction	3
1.1. Questions for reviewers	4
2. Terminology	4
3. Inter-domain SMTP with TLSA	4
3.1. MX lookup checks	5
3.2. SMTP server checks	5
4. Intra-domain SMTP with TLSA	6
5. The Transmitted: header field	7
6. IANA Considerations	8
6.1. "with" protocol types	8
6.2. Permanent message header field registration	8
6.3. "dane" MTA-name-type	8
7. Security considerations	9
7.1. Fallback to insecure SMTP	9
7.2. A mail domain trusts its SMTP servers	9
7.3. Temporary failures and denial of service	9
7.4. Deliberate omissions	10
8. Internationalization Considerations	10
9. Acknowledgements	10
10. References	10
10.1. Normative References	10
10.2. Informative References	12
Appendix A. Example	12
Appendix B. Rationale - choice of certificate identity	13
Appendix C. Change log	13
C.1. Changes in version -04	13
C.2. Changes in version -03	14
C.3. Changes in version -02	14
C.4. Changes in version -01	14
Author's Address	14

1. Introduction

The specification for SMTP over TLS [RFC3207] does not describe how to authenticate a server: which identity relating to the connection ought to be authenticated by the server's certificate. In practice, most certificates presented by publicly-referenced SMTP servers either cannot be validated with respect to a well-known certification authority, or do not verify any identity expected by the client.

As a result, inter-domain SMTP clients cannot require working server authentication if they want to successfully send mail using TLS. Therefore TLS currently provides only a limited amount of additional security for inter-domain SMTP. Its encryption protects against on-path passive eavesdropping; but it does not protect against an active attack, since the client has no way to detect when an attacker is spoofing the server.

This memo describes how to fix this using DNSSEC [RFC4033] and TLSA records [I-D.ietf-dane-protocol] with owner names of the form "_25._tcp.hostname".

We use DNSSEC to secure the association between a mail domain and its SMTP server host names, and between the host names and their certificates. Connections to servers are authenticated by their TLS certificates.

As well as its normal function of providing an association between a domain name and a certificate, we are also using the existence of a TLSA record to signal to the client that it can expect the server to offer TLS with a valid certificate.

The security situation is better for intra-domain SMTP, because in this case the client and server can be configured with prior knowledge of how to authenticate each other. This specification can also be used for authenticating servers in intra-domain SMTP.

This memo does not cover message submission [RFC4409] [RFC5068] [RFC6186], nor does it cover LMTP [RFC2033], since they use the DNS in a different way than MTA-to-MTA SMTP.

The protocol described in this memo adds new security checks that can cause email delivery to be delayed when a security failure is detected. We specify that clients treat a problems as a "temporary failure", causing the message to be queued for a later delivery attempt, in the hope that the attack (or configuration error) will have been dealt with.

1.1. Questions for reviewers

Is the Transmitted: header useful enough to include in this spec?
Should it be dropped, or perhaps moved to another document?

Is the "dane" MTA-name-type for use in delivery status notifications
a good idea? Is it likely to cause interoperability problems?

Is the description of DNSSEC validation over-done? Can it be trimmed
down so it relies more on the core DNSSEC specs?

2. Terminology

ADMD: An ADministrative Management Domain, as described in the
Internet Mail Architecture [RFC5598].

Inter-domain SMTP: SMTP between different ADMDs across the public
Internet, where a client MTA sends mail to a publicly-referenced
SMTP server MTA.

Intra-domain SMTP: SMTP between MTAs within an ADMD.

Mail domain: The part of an email address after the "@"; also the
owner name of a (possibly implicit) MX record.

MX resolution: The algorithm for resolving a mail domain into a set
of SMTP server hosts, described in [RFC5321] section 5.

Publicly-referenced SMTP server: An SMTP server which runs on port
25 of an Internet host located using MX resolution. (This term is
from [RFC3207].)

SMTP server host name: The target of a (possibly implicit) MX
record.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
memo are to be interpreted as described in [RFC2119].

3. Inter-domain SMTP with TLSA

In the following we describe some additions to the usual MX
resolution algorithm described in [RFC5321] section 5. If there is
any conflict between this memo and the other specifications cited in
this section, that is an error in this memo.

3.1. MX lookup checks

The client SHALL look up the MX RRset for the mail domain. There are three succesful results that yield a list of SMTP server host names:

- o A list of one or more MX records;
- o An implicit MX record, in lieu of an empty list of MX records;
- o A CNAME or DNAME pointing to a successful result.

If the lookup is not successful, the client SHALL proceed as described in [RFC5321] section 5.

If any of the responses is "bogus" according to DNSSEC validation ([RFC4033] section 5) the client MUST treat this as a temporary error.

For this protocol to take effect, all of these DNS RRsets MUST be "secure" according to DNSSEC validation. In the case of an implicit MX record, there MUST be a secure denial of existence of an MX RRset for the mail domain. In the case of a (chain of) CNAME or DNAME RRs, the whole chain MUST be secure as well as the ultimate target.

If they are not all secure, this protocol has not been fully deployed. The client SHOULD fall back to insecure delivery (which might be over unauthenticated TLS).

(If the client is using a non-validating security-aware stub resolver (see [RFC4033] section 7), it can rely on its recursive name server to perform these checks and set the AD bit according to the result - see [RFC4035] section 3.2.3.)

The client now has an authentic list of SMTP server host names and priority values. It processes this list as described in [RFC5321] section 5 (sorting the host names etc.) without regard to the presence or absence of DNSSEC or TLSA records.

3.2. SMTP server checks

This sub-section applies to each SMTP server host name individually.

When connecting to a server, the client SHALL look up the server's TLSA RRset as described in [I-D.ietf-dane-protocol] section 3. That is, the TLSA RRset owner name SHALL be "_25._tcp.hostname" where "hostname" is the SMTP server host name. The response can be one of the following (as listed in [I-D.ietf-dane-protocol] section 4.1):

- o A secure answer containing one or more TLSA records, in which case the client SHALL proceed as described below.
- o A bogus answer or other failure, which the client MUST treat as a temporary error.
- o If there is no TLSA record or its DNSSEC validation state is insecure or indeterminate, this protocol has not been fully deployed. The client SHOULD deliver to this server insecurely (which might be over unauthenticated TLS).

The client now has one or more TLSA records for the server it is connecting to.

The client MUST ensure that the server offers the STARTTLS service extension [RFC3207] in its response to the client's EHLO command ([RFC5321] section 4.1.1.1).

The client SHALL then issue the STARTTLS command which MUST be successful. It then proceeds with TLS negotiation [RFC5246]. If the client uses the Server Name Indication TLS extension ([RFC6066] section 3) it MUST use the SMTP server host name as the value for the ServerName field.

The client SHALL validate the server's certificate as described in [I-D.ietf-dane-protocol] section 2.1.

The client SHALL verify the server's identity as described in [RFC6125] section 6. Its list of reference identifiers SHOULD include the SMTP server host name with type DNS-ID, and MAY include a second copy of the host name with type CN-ID.

If any of these checks fail, the client MUST disconnect from the server and treat this as a temporary failure.

The client can now proceed to deliver mail securely.

4. Intra-domain SMTP with TLSA

Mail transmission within an ADMD can be based on MX records (such as when delivering incoming mail to its destination host) or on statically configured host names (such as when routing outgoing mail via a border relay).

When routing internal mail using MX records, Section 3 applies the same as for inter-domain SMTP.

When routing mail using host names, the MX lookup step is skipped and only Section 3.2 applies.

5. The Transmitted: header field

The client MAY wish to insert a Transmitted: header field at the start of the message header just before transmitting the message. This records the result of the checks specified in the previous section. (See Section 7 for some comments on its utility or lack thereof.) It is a client-side counterpart to the Received: header field ([RFC5321] section 4.4) and has very similar syntax. It SHOULD be treated as a trace field.

The syntax of the Transmitted: header field is described using ABNF [RFC5234]. Non-terminal syntax rules not defined in this memo are defined in [RFC5321], or [RFC5322], or [RFC5234].

```
Transmitted-line = "Transmitted:" FWS To-domain By-domain
                  Opt-info [CFWS] ";" date-time CRLF
```

```
To-domain       = "TO" FWS Extended-Domain
```

A <Transmitted-line> SHALL include:

- o A <To-domain> clause describing the SMTP server. The <Domain> part of a <To-domain> SHALL be the same as the SMTP server host name.
- o A <By-domain> clause identifying the SMTP client that added the header. (If the client also acts as a server this is the same <By-domain> clause it would include in any Received: header fields it adds.) This clause helps with recovery if the original order of a message header's fields has been lost.
- o Various <Opt-info> clauses, which MUST include a <With> clause. The <Protocol> part of this clause is used to indicate whether the client successfully authenticated the server, using one of the types specified in Section 6.1.
- o And a <date-time> to further help with disordering in case a message is transmitted by the same client more than once.

6. IANA Considerations

6.1. "with" protocol types

The "with" protocol type registry includes a number of keywords that indicate the use of SMTP with or without TLS and/or AUTH [RFC3848]. When these types appear in a Transmitted: header field "with" clause they indicate that the client did not authenticate the server as described in Section 3.

- o The new keyword "ESMTPT" indicates the use of ESMTP [RFC5321] with STARTTLS [RFC3207] when the client successfully authenticated the server.
- o The new keyword "ESMTPTA" indicates the use of ESMTP [RFC5321] with STARTTLS [RFC3207] and AUTH [RFC4954] when the client successfully authenticated the server.

These new keywords are not for use in Received: header fields since the server cannot tell whether or not the client authenticated it.

There are no keywords corresponding to a client trying and failing to authenticate the server, since in this case no message transmission occurs.

6.2. Permanent message header field registration

Header field name: Transmitted:

Applicable protocol: mail

Status: standard

Change controller: IETF

Specification document this memo

6.3. "dane" MTA-name-type

Delivery status notifications [RFC3464] can include a Remote-MTA field recording an SMTP server host name. When this has been authenticated according to Section 3 the reporting MTA MAY use an MTA-type-name of "dane".

- a. MTA-type-name: "dane"
- b. Syntax: same as the "dns" MTA-type-name [RFC3461]

- c. Translation into US-ASCII: none needed

7. Security considerations

7.1. Fallback to insecure SMTP

This memo provides only conditional security. It allows a server to publish in the DNS the details of how it can be authenticated. Clients that implement this protocol can use it to provide a strong guarantee that they are sending mail to the correct place. If either of these is missing, mail delivery will be insecure.

There is no secure way for a server to tell if a client has authenticated it using this protocol. This is a general limitation of TLS. The Transmitted: header field records this information for tracing and debugging and measuring deployment, not for security purposes.

We do not specify that clients check that all of a mail domain's SMTP server host names are consistent in whether they have or do not have TLSA records. This is so that partial or incremental deployment does not break mail delivery. Different levels of deployment are likely if a domain has a third-party backup MX, for example.

The MX sorting rules are unchanged; in particular they have not been altered in order to prioritize secure servers over insecure servers. If a site wants to be secure it needs to deploy this protocol completely; a partial deployment is not secure and we make no special effort to support it.

7.2. A mail domain trusts its SMTP servers

By signing their zone with DNSSEC, a mail domain owner implicitly instructs SMTP clients to check their SMTP server TLSA records. This implies another point in the trust relationship between mail domain owner and smtp server operator. Most of the setup requirements for this protocol fall on the SMTP server operator: installing a TLS certificate with the correct name, and publishing a TLSA record under that name. If these are not correct then mail delivery from TLSA-aware clients might be delayed.

7.3. Temporary failures and denial of service

Many provisioning failures in SMTP cause "permanent" failures, that is the immediate and final rejection of the message. This includes missing DNS records, an SMTP server that is not configured to accept mail for the recipient domain, and so forth.

In this protocol, provisioning an incorrect TLS certificate triggers a temporary error. This is because we want to minimise the damage that occurs when an on-path attacker intercepts the TCP connection between an SMTP client and server. An attacker can cause delays, but is not able to trigger immediate delivery failures.

7.4. Deliberate omissions

We do not specify that clients check the DNSSEC state of the SMTP server address records. This is not necessary since the certificate checks ensure that the client has connected to the correct server. (The address records will normally have the same security state as the TLSA records, but they can differ if there are CNAME or DNAME indirections.)

This memo does not specify any changes to SMTP client authentication. Inter-domain SMTP client authentication remains extremely weak. Intra-domain SMTP can be configured as strong as necessary (using SMTP AUTH or TLS client certificates, for instance) but that is out of scope for this memo.

8. Internationalization Considerations

If any of the DNS queries are for an internationalized domain name, then they need to use the A-label form [RFC5890].

9. Acknowledgements

Thanks to Mark Andrews for arguing that authenticating the SMTP server host name is the right thing, and that we ought to rely on DNSSEC to secure the MX lookup. Thanks to Ned Freed, Olafur Gudmundsson, Paul Hoffman, Phil Pennock, Hector Santos, and Alessandro Vesely for helpful suggestions.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3207] Hoffman, P., "SMTP Service Extension for Secure SMTP over Transport Layer Security", RFC 3207, February 2002.
- [RFC3461] Moore, K., "Simple Mail Transfer Protocol (SMTP) Service

- Extension for Delivery Status Notifications (DSNs)", RFC 3461, January 2003.
- [RFC3464] Moore, K. and G. Vaudreuil, "An Extensible Message Format for Delivery Status Notifications", RFC 3464, January 2003.
- [RFC3848] Newman, C., "ESMTP and LMTP Transmission Types Registration", RFC 3848, July 2004.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.
- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.
- [RFC4954] Siemborski, R. and A. Melnikov, "SMTP Service Extension for Authentication", RFC 4954, July 2007.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.
- [RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, October 2008.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, August 2010.
- [RFC6066] Eastlake, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, January 2011.
- [RFC6125] Saint-Andre, P. and J. Hodges, "Representation and Verification of Domain-Based Application Service Identity within Internet Public Key Infrastructure Using X.509 (PKIX) Certificates in the Context of Transport Layer Security (TLS)", RFC 6125, March 2011.
- [I-D.ietf-dane-protocol]
Hoffman, P. and J. Schlyter, "The DNS-Based Authentication

of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", draft-ietf-dane-protocol-23 (work in progress), May 2012.

10.2. Informative References

- [RFC2033] Myers, J., "Local Mail Transfer Protocol", RFC 2033, October 1996.
- [RFC4409] Gellens, R. and J. Klensin, "Message Submission for Mail", RFC 4409, April 2006.
- [RFC5068] Hutzler, C., Crocker, D., Resnick, P., Allman, E., and T. Finch, "Email Submission Operations: Access and Accountability Requirements", BCP 134, RFC 5068, November 2007.
- [RFC5598] Crocker, D., "Internet Mail Architecture", RFC 5598, July 2009.
- [RFC6186] Daboo, C., "Use of SRV Records for Locating Email Submission/Access Services", RFC 6186, March 2011.

Appendix A. Example

In the following, most of the DNS resource data is elided for simplicity.

```
; mail domain
example.com.      MX      1 mx.example.net.
example.com.      RRSIG   MX ...

; SMTP server host name
mx.example.net.   A        192.0.2.1
mx.example.net.   AAAA     2001:db8:212:8::e:1

; TLSA resource record
_25._tcp.mx.example.net.  TLSA   ...
_25._tcp.mx.example.net.  RRSIG  TLSA ...
```

Mail for addresses at example.com is delivered by SMTP to mx.example.net. Connections to mx.example.net port 25 that use STARTTLS will get a server certificate that authenticates the name mx.example.net.

Appendix B. Rationale - choice of certificate identity

There are a number of reasons for the certificate to authenticate the SMTP server host name rather than the mail domain.

SMTP allows a client to transfer mail to recipients at multiple domains in the same connection. If the certificate identifies the host name then it does not need to list all the possible mail domains.

It is not in general feasible for the server to select a mail domain certificate based on the recipient domains when the connection is established (using Server Name Indication, [RFC6066] section 3), because an SMTP client might not know all of the recipients when it establishes the connection.

Outgoing SMTP relays and message submission servers handle mail for any domain, so in those cases the only sensible option is for the certificate to contain the host name. It is more consistent for incoming MX server certificates to match.

It is common for SMTP servers to act in multiple roles, as outgoing relays or as incoming MX servers, depending on the client identity. It is simpler if the server can present the same certificate regardless of the role in which it is to act.

Sometimes the server does not know its role until the client has authenticated, which usually occurs after TLS has been established.

This protocol does not provide an option for directly authenticating the mail domain because that would add complexity without providing any benefit, and security protocols are best kept simple. As described above, there are real-world cases where authenticating the mail domain cannot be made to work, so there are complicated criteria for when mail domain TLSA records might be used and when they cannot. This is all avoided by authenticating the SMTP server host name.

Finally, this protocol only affects the logic in the SMTP client and requires no additional SMTP server functionality, such as support for the TLS Server Name Indication extension.

Appendix C. Change log

C.1. Changes in version -04

Add some questions for reviewers

Add a note about stub resolvers and the AD bit.

Internationalization considerations.

C.2. Changes in version -03

Clarify how to use SNI with this protocol.

Clarify lack of changes to MX sorting rules.

Mention DNAME as well as CNAME.

An example.

C.3. Changes in version -02

Clarify the wording that describes how a client determines that this protocol is in effect.

Divide the security considerations into sub-sections, and add a subsection on denial of service.

Clarify intro, mentioning TLSA owner name format.

Extend the scope to cover MTA-to-MTA mail within an ADMD as well as between ADMDs.

C.4. Changes in version -01

More about why not to authenticate mail domains in the rationale.

Change DNS-ID requirement from MUST to SHOULD to follow RFC 6125.

Acknowledgments section.

Transmitted: header trace field. Not sure if this is a good idea; feedback wanted.

"dane" MTA-name-type for use in DSNs. Even less sure if this is a good idea.

Author's Address

Tony Finch
University of Cambridge Computing Service
New Museums Site
Pembroke Street
Cambridge CB2 3QH
ENGLAND

Phone: +44 797 040 1426
Email: dot@dotat.at
URI: <http://dotat.at/>

