

Internet Draft  
[draft-hoffman-smtp-ssl-10.txt](#)  
November 6, 1998  
Expires in six months

Paul Hoffman  
Internet Mail Consortium

## SMTP Service Extension for Secure SMTP over TLS

Status of this memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

To learn the current status of any Internet-Draft, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ftp.ietf.org (US East Coast), or ftp.isi.edu (US West Coast).

### **1. Abstract**

This document describes an extension to the SMTP service that allows an SMTP server and client to use transport-layer security to provide private, authenticated communication over the Internet. This gives SMTP agents the ability to protect some or all of their communications from eavesdroppers and attackers.

### **2. Introduction**

SMTP [[RFC-821](#)] servers and clients normally communicate in the clear over the Internet. In many cases, this communication goes through one or more router that is not controlled or trusted by either entity. Such an untrusted router might allow a third party to monitor or alter the communications between the server and client.

Further, there is often a desire for two SMTP agents to be able to authenticate each others' identities. For example, a secure SMTP server might only allow communications from other SMTP agents it knows, or it might act differently for messages received from an agent it knows than from one it doesn't know.

TLS [[TLS](#)], more commonly known as SSL, is a popular mechanism for enhancing TCP communications with privacy and authentication. TLS is in wide use with the HTTP protocol, and is also being used for adding security to many other

common protocols that run over TCP.

## **2.1 Terminology**

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC-2119](#)].

## **3. STARTTLS Extension**

The STARTTLS extension to SMTP is laid out as follows:

- (1) the name of the SMTP service defined here is STARTTLS;
- (2) the EHLO keyword value associated with the extension is STARTTLS;
- (3) the STARTTLS keyword has no parameters;
- (4) a new SMTP verb, "STARTTLS", is defined;
- (5) no additional parameters are added to any SMTP command.

## **4. The STARTTLS Keyword**

The STARTTLS keyword is used to tell the SMTP client that the SMTP server allows use of TLS. It takes no parameters.

## **5. The STARTTLS Command**

The format for the STARTTLS command is:

STARTTLS

with no parameters.

After the client gives the STARTTLS command, the server responds with one of the following reply codes:

**220 Ready to start TLS**

**501 Syntax error (no parameters allowed)**

**454 TLS not available due to temporary reason**

A publicly-referenced SMTP server MUST NOT require use of the STARTTLS extension in order to deliver mail locally. This rule prevents the STARTTLS extension from damaging the interoperability of the Internet's SMTP infrastructure. A publicly-referenced SMTP server is an SMTP server which runs on port 25 of an Internet host listed in the MX record (or A record if an MX record is not present) for the domain name on the right hand side of an Internet mail address.

Any SMTP server may refuse to accept messages for relay based on authentication supplied during the TLS negotiation. An SMTP server that is not publicly referenced may refuse to accept any messages for relay or local delivery based on authentication supplied during the TLS negotiation.

A SMTP server that is not publicly referenced may choose to require that the client perform a TLS negotiation before accepting any commands. In this case, the server SHOULD return the reply code:

#### **530 Must issue a STARTTLS command first**

to every command other than NOOP, EHLO, STARTTLS, or QUIT. If the client and server are using the ENHANCEDSTATUSCODES ESMTP extension [[RFC-2034](#)], the status code to be returned SHOULD be 5.7.0.

After receiving a 220 response to a STARTTLS command, the client SHOULD start the TLS negotiation before giving any other SMTP commands.

If the SMTP client is using pipelining as defined in [RFC 1854](#), the STARTTLS command must be the last command in a group.

#### **5.1 Processing After the STARTTLS Command**

After the TLS handshake has been completed, both parties MUST immediately decide whether or not to continue based on the authentication and privacy achieved. The SMTP client and server may decide to move ahead even if the TLS negotiation ended with no authentication and/or no privacy because most SMTP services are performed with no authentication and no privacy, but some SMTP clients or servers may want to continue only if a particular level of authentication and/or privacy was achieved.

If the SMTP client decides that the level of authentication or privacy is not high enough for it to continue, it SHOULD issue an SMTP QUIT command immediately after the TLS negotiation is complete. If the SMTP server decides that the level of authentication or privacy is not high enough for it to continue, it SHOULD reply to every SMTP command from the client (other than a QUIT command) with the 554 reply code (with a possible text string such as "Command refused due to lack of security").

The decision of whether or not to believe the authenticity of the other party in a TLS negotiation is a local matter. However, some general rules for the decisions are:

- A SMTP client would probably only want to authenticate an SMTP server whose server certificate has a domain name that is the domain name that the client thought it was connecting to.
- A publicly-referenced SMTP server would probably want to accept any certificate from an SMTP client, and would possibly want to put distinguishing information about the certificate in the Received header of messages that were relayed or submitted from the client.

#### **5.2 Result of the STARTTLS Command**

Upon completion of the TLS handshake, the SMTP protocol is reset to the initial state (the state in SMTP after a server issues a 220 service ready greeting). The server MUST discard any knowledge obtained from the client, such as the argument to the EHLO command, which was not obtained from the TLS negotiation itself. The client MUST discard any knowledge obtained from the server, such as the list of SMTP service extensions, which was not obtained from the TLS negotiation itself. The client SHOULD send an EHLO command as the first command after a successful TLS negotiation.

The list of SMTP service extensions returned in response to an EHLO command received after the TLS handshake MAY be different than the list returned before the TLS handshake. For example, an SMTP server might not want to advertise support for a particular SASL mechanism [[SASL](#)] unless a client has sent an appropriate client certificate during a TLS handshake.

Both the client and the server MUST know if there is a TLS session active. A client MUST NOT attempt to start a TLS session if a TLS session is already active. A server MUST NOT return the TLS extension in response to an EHLO command received after a TLS handshake has completed.

## **6. Usage Example**

The following dialog illustrates how a client and server can start a TLS session:

```
S: <waits for connection on TCP port 25>
C: <opens connection>
S: 220 mail.imc.org SMTP service ready
C: EHLO mail.ietf.org
S: 250-mail.imc.org offers a warm hug of welcome
S: 250 STARTTLS
C: STARTTLS
S: 220 Go ahead
C: <starts TLS negotiation>
C & S: <negotiate a TLS session>
C & S: <check result of negotiation>
C: <continues by sending an SMTP command>
. . .
```

## **7. Security Considerations**

It should be noted that SMTP is not an end-to-end mechanism. Thus, if an SMTP client/server pair decide to add TLS privacy, they are not securing the transport from the originating mail user agent to the recipient. Further, because delivery of a single piece of mail may go between more than two SMTP servers, adding TLS privacy to one pair of servers does not mean that the entire SMTP chain has been made private. Further, just because an SMTP server can authenticate an SMTP client, it does not mean that the mail from the SMTP client was authenticated by the SMTP client

when the client received it.

Both the SMTP client and server must check the result of the TLS negotiation to see whether acceptable authentication or privacy was achieved. Ignoring this step completely invalidates using TLS for security. The decision about whether acceptable authentication or privacy was achieved is made locally, is implementation-dependant, and is beyond the scope of this document.

The SMTP client and server should note carefully the result of the TLS negotiation. If the negotiation results in no privacy, or if it results in privacy using algorithms or key lengths that are deemed not strong enough, or if the authentication is not good enough for either party, the client may choose to end the SMTP session with an immediate QUIT command, or the server may choose to not accept any more SMTP commands.

A server announcing in an EHLO response that it uses a particular TLS protocol should not pose any security issues, since any use of TLS will be at least as secure as no use of TLS.

A man-in-the-middle attack can be launched by deleting the "250 STARTTLS" response from the server. This would cause the client not to try to start a TLS session. An SMTP client can protect against this attack by recording the fact that a particular SMTP server offers TLS during one session and generating an alarm if it does not appear in the EHLO response for a later session. The lack of TLS during a session SHOULD NOT result in the bouncing of email, although it could result in delayed processing.

Before the TLS handshake has begun, any protocol interactions are performed in the clear and may be modified by an active attacker. For this reason, clients and servers MUST discard any knowledge obtained prior to the start of the TLS handshake upon completion of the TLS handshake.

The STARTTLS extension is not suitable for authenticating the author of an email message unless every hop in the delivery chain, including the submission to the first SMTP server, is authenticated. Another proposal [[SMTP-AUTH](#)] can be used to authenticate delivery and MIME security multiparts [MIME-SEC] can be used to authenticate the author of an email message. In addition, the [[SMTP-AUTH](#)] proposal offers simpler and more flexible options to authenticate an SMTP client and the SASL EXTERNAL mechanism [[SASL](#)] MAY be used in conjunction with the STARTTLS command to provide an authorization identity.

## **[A. References](#)**

[RFC-821] "Simple Mail Transfer Protocol", [RFC 821](#)

[RFC-1869] "SMTP Service Extensions", [RFC 1869](#)

[RFC-2034] "SMTP Service Extension for Returning Enhanced Error Codes", RFC 2034

[RFC-2119] "Key words for use in RFCs to Indicate Requirement Levels",  
[BCP 14](#), [RFC 2119](#)

[SASL] "Simple Authentication and Security Layer (SASL)", [RFC 2222](#)

[SMTP-AUTH] "SMTP Service Extension for Authentication",  
Internet Draft [draft-myers-smtp-auth-xx.txt](#)

[TLS] "The TLS Protocol Version 1.0", [draft-ietf-tls-protocol-xx.txt](#)

## **[B](#). Author's Address**

Paul Hoffman  
Internet Mail Consortium  
[127](#) Segre Place  
Santa Cruz, CA 95060  
(408) 426-9827  
[phoffman@imc.org](mailto:phoffman@imc.org)