

Intended status: Standards Track
Expires: November 2008

MTLS: TLS Multiplexing
draft-badra-hajjeh-mtls-04.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on November 19, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

The Transport Layer Security (TLS) standard provides connection security with mutual authentication, data confidentiality and integrity, key generation and distribution, and security parameters

negotiation. However, missing from the protocol is a way to multiplex application data over a single TLS session.

This document defines MTLS, an application-level protocol running over TLS (or DTLS) Record protocol. The MTLS design provides application multiplexing over a single TLS (or DTLS) session. Therefore, instead of associating a TLS connection with each application, MTLS allows several applications to protect their exchanges over a single TLS session.

Table of Contents

| | |
|--|-------------------|
| 1. Introduction..... | 2 |
| 1.1. Conventions used in this document..... | 3 |
| 2. TLS multiplexing overview and considerations..... | 3 |
| 2.1. Handshake..... | 3 |
| 2.1.1. Opening and closing connections..... | 4 |
| 2.2. MTLS protocol..... | 5 |
| 2.3. MTLS Message Types..... | 7 |
| 3. Security Considerations..... | 7 |
| 4. IANA Considerations..... | 8 |
| 5. References..... | 8 |
| 5.1. Normative References..... | 8 |
| 5.2. Informative References..... | 8 |
| Author's Addresses..... | 9 |
| Intellectual Property and Copyright Statements..... | 9 |

[1. Introduction](#)

HTTP over TLS [[RFC2818](#)], POP over TLS and IMAP over TLS [[RFC2595](#)] are examples of securing, respectively HTTP, POP and IMAP data exchanges using the TLS protocol [[I-D.ietf-tls-rfc4346-bis](#)].

TLS ([[I-D.ietf-tls-rfc4346-bis](#)], [[RFC4347](#)]) is the most deployed security protocol for securing exchanges, for authenticating entities and for generating and distributing cryptographic keys. However, what is missing from the protocol is the way to multiplex application data over the same TLS session.

Actually, TLS (or DTLS) clients and servers MUST establish a TLS (or DTLS) session for each application they want to run over a transport layer. However, some applications may agree or be configured to use the same security policies or parameters (e.g. authentication method and cipher_suite) and then to share a single TLS session to protect their exchanges. In this way, this document describes a way to allow application multiplexing over TLS.

The document motivations included:

- o TLS is application protocol-independent. Higher-level protocol can operate on top of the TLS protocol transparently.
- o TLS is a protocol of a modular nature. Since TLS is developed in four independent protocols, the approach defined in this document can be used with a total reuse of pre-existing TLS infrastructures and implementations.
- o It provides a secure VPN tunnel over a transport layer. Unlike "ssh-connection" [[RFC4254](#)], MTLS can run over unreliable transport protocols, such as UDP.
- o Establishing a single TLS session for a number of applications - instead of establishing a TLS session per one of those applications- reduces resource consumption, latency and messages flow that are associated with executing simultaneous TLS sessions.
- o TLS can not forbid an intruder to analyze the traffic and cannot protect data from inference. Thus, the intruder can know the type of application data transmitted through the TLS session. However, the approach defined in this document allows, by its design, data protection against inference.

[1.1](#). Conventions used in this document

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](#) [[RFC2119](#)].

[2](#). TLS multiplexing overview and considerations

This document defines an application-level protocol called Multiplexing TLS (MTLS) to handle data multiplexing.

[2.1](#). Handshake

If the client is willing to run MTLS, it MUST connect to the server that passively listens for the incoming TLS connection on the IANA-to-be-assigned TCP or UDP port <TBA>. The client MUST therefore send the TLS ClientHello to begin the TLS handshake. Once the Handshake is complete, the client and the server can establish and manage many application channels using the MTLS requests/responses defined below.

2.1.1. Opening and closing connections

Once the Handshake is complete, both the client and the server can start data multiplexing using a set of requests/responses defined below. All requests/responses will pass through MTLS layer and are formatted into MTLS packets, depending on each request/response.

The sender MAY request the opening of many channels. For each channel, the MTLS layer generates and sends the following request:

```
struct {  
    uint8  type;  
    opaque sender_channel_id[2];  
    uint32 sender_window_length;  
    uint32 sender_max_packet_length;  
    opaque source_address_machine<4..7>;  
    opaque source_port[2];  
    opaque destination_address_machine<4..7>;  
    opaque destination_port[2];  
} RequestEstablishmentChannel;
```

The field "type" specifies the MTLS packet type (types are summarized below), the "max_packet_length" and the "sender_channel_id" are used as described below. The "source_address_machine" MAY carry either the numeric IP address or the domain name of the host from where the application originates the data multiplexing request and the "port" is the port on the host from where the connection originated.

The sender initializes its "window_length" with the data length (in octets), specifying how many bytes the receiver can maximally send on the channel before receiving a new window length (available free space). Each end of the channel establishes a "receive buffer" and a "send buffer".

The sender initializes its "max_packet_length" with the data length (in octets), specifying the maximal packet's length in octets the receiver can send on the channel.

The "destination_address_machine" and "destination_port" specify the TCP/IP host and port where the recipient should connect the channel. The "destination_address_machine" MAY be either a domain name or a numeric IP address.

The receiver decides whether it can open the channel, and replies with one of the following messages:


```
struct {
    uint8  type;
    opaque sender_channel_id[2];
    opaque receiver_channel_id[2];
    uint32 receiver_window_length;
    uint32 max_packet_length;
} RequestEstablishmentSuccess;

struct {
    uint8  type;
    opaque sender_channel_id[2];
    opaque error<0..2^16>;
} RequestEchecChannel;
```

The field "error" conveys a description of the error.

If an error occurs at the MTLS layer, the established secure session is still valid and no alert of any type is sent by the TLS Record.

Each MTLS channel has its identifier computed as:

$$\text{channel_id} = \text{sender_channel_id} + \text{"receiver_channel_id"}$$

Where "+" indicates concatenation.

The following packet MAY be sent to notify the receiver that the sender will not send any more data on this channel and that any data received after a closure request will be ignored. The sender of the closure request MAY close its "receive buffer" without waiting for the receiver's response. However, the receiver MUST respond with a confirmation of the closure and close down the channel immediately, discarding any pending writes.

```
struct {
    uint8  type;
    opaque channel_id[4];
} CloseChannel;

struct {
    uint8  type;
    opaque channel_id[4];
} ConfirmationCloseChannel;
```

[2.2. MTLS protocol](#)

The structure of the MTLS packet is described below. The "sender_channel_id" and "receiver_channel_id" are the same generated

during the connection establishment. The length conveys the data length of the current packet.

Each entity maintains its "max_packet_length" (that is originally initialized during the connection establishment) to a value not bigger than the maximum size of this entity's "receive buffer". For each received packet, the entity MUST subtract the packet's length from the "max_packet_length". The result is always positive since the packet's length is always less than or equal to the current "max_packet_length".

The free space of the "receive buffer" MAY increase in length. Consequently, the entity MUST inform the other end about this increase, allowing the other entity to send packet with length bigger than the old "max_packet_length" but smaller or equal than the new value.

The entity MAY indicate this increase by sending an Acknowledgment packet. The Acknowledgment packet carries the available free space ("free_space" field in octets) the receiver of that packet can send on the channel before receiving a new window length.

If the length of the "receive buffer" does not change, Acknowledgment packet will never be sent.

In the case where the "receive buffer" of an entity fills up, the other entity MUST wait for an Acknowledgment packet before sending any more MTLSPlaintext packets.

```
struct {
    uint8  type;
    opaque channel_id[4];
    uint32 length;
    opaque data[MTLSPlaintext.length];
} MTLSPlaintext;

struct {
    uint8  type;
    opaque channel_id[4];
    uint32 free_space;
} Acknowledgment;
```

The TLS Record Layer receives data from MTLS, supposes it as uninterpreted data and applies the fragmentation and the cryptographic operations on it, as defined in [I-D.ietf-tls-rfc4346-bis]. The type is set to mtls(TBA).

Note: multiple MTLS fragments MAY be coalesced into a single TLSPlaintext record.

Received data is decrypted, verified, decompressed, and reassembled, then delivered to MTLS entity. Next, the MTLS sends data to the appropriate application using the channel identifier and the length value.

2.3. MTLS Message Types

This section defines the initial set of MTLS Message Types used in Request/Response exchanges. The Message Type field is one octet and identifies the structure of an MTLS Request or Response message.

The messages defined in this document are listed below. More Message Types may be defined in future documents. The list of Message Types, as defined through this document, is maintained by the Internet Assigned Numbers Authority (IANA). Thus, an application needs to be made to the IANA in order to obtain a new Message Type value. Since there are subtle (and not-so-subtle) interactions that may occur in this protocol between new features and existing features that may result in a significant reduction in overall security, new values SHALL be defined only through the IETF Consensus process specified in [\[RFC2434\]](#). All of the messages defined in this document follow the convention that for each message that an entity sends and that the other entity understands, this latter entity replies with a message of the same type.

| | |
|-----------------------------|---|
| RequestEstablishmentChannel | 1 |
| RequestEstablishmentSuccess | 2 |
| RequestEchecChannel | 3 |
| CloseChannel | 4 |
| ConfirmationCloseChannel | 5 |
| MTLSPlaintext | 6 |
| Acknowledgment | 7 |

3. Security Considerations

Security issues are discussed throughout this document and in [I-D.ietf-tls-rfc4346-bis] and [\[RFC4347\]](#) documents.

If a fatal error related to any channel or a connection of an arbitrary application occurs, the secure session MUST NOT be resumed. This is logic since the Record protocol does not distinguish between the MTLS channels. However, if an error occurs at the MTLS layer, both parties immediately close the related channels, but not the TLS session (no alert of any type is sent by the TLS Record).

4. IANA Considerations

This section provides guidance to the IANA regarding registration of values related to the TLS protocol.

IANA is requested to assign a TCP and UDP port numbers that will be the default port for MTLS sessions as defined in this document.

There is one name space in MTLS that requires registration: Message Types.

Message Types have a range from 1 to 255, of which 1-7 are to be allocated for this document. Because a new Message Type has considerable impact on interoperability, a new Message Type SHALL be defined only through the IETF Consensus process specified in [\[RFC2434\]](#).

5. References

5.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 2434](#), October 1998.
- [I-D.ietf-tls-rfc4346-bis]
Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [draft-ietf-tls-rfc4346-bis-10](#) (work in progress), March 2008.
- [RFC4347] Rescorla, E., Modadugu, N., "Datagram Transport Layer Security", [RFC 4347](#), April 2006.

5.2. Informative References

- [RFC2818] Rescorla, E., "HTTP Over TLS", [RFC 2818](#), May 2000.
- [RFC2595] Newman, C., "Using TLS with IMAP, POP3 and ACAP", [RFC 2595](#), June 1999.
- [RFC4254] Lonvick, C., "SSH Connection Protocol", [RFC 4254](#), January 2005.

Author's Addresses

Mohamad Badra
LIMOS Laboratory - UMR6158, CNRS
France

Email: badra@isima.fr

Ibrahim Hajjeh
INEOVATION
France

Email: hajjeh@ineovation.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF

THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.