

TLS  
Internet-Draft  
Obsoletes: [6347](#) (if approved)  
Intended status: Standards Track  
Expires: April 29, 2017

E. Rescorla  
RTFM, Inc.  
H. Tschofenig  
ARM Limited  
October 26, 2016

**The Datagram Transport Layer Security (DTLS) Protocol Version 1.3**  
**draft-rescorla-tls-dtls13-00**

Abstract

This document specifies Version 1.3 of the Datagram Transport Layer Security (DTLS) protocol. DTLS 1.3 allows client/server applications to communicate over the Internet in a way that is designed to prevent eavesdropping, tampering, and message forgery.

The DTLS 1.3 protocol is intentionally based on the Transport Layer Security (TLS) 1.3 protocol and provides equivalent security guarantees. Datagram semantics of the underlying transport are preserved by the DTLS protocol.

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 April 29, 2017.

Copyright Notice

Copyright (c) 2016 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">3</a>
<a href="#">2.</a>	<a href="#">Conventions and Terminology</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">DTLS Design Rational and Overview</a>	<a href="#">4</a>
<a href="#">3.1.</a>	<a href="#">Packet Loss</a>	<a href="#">5</a>
<a href="#">3.1.1.</a>	<a href="#">Reordering</a>	<a href="#">6</a>
<a href="#">3.1.2.</a>	<a href="#">Message Size</a>	<a href="#">6</a>
<a href="#">3.2.</a>	<a href="#">Replay Detection</a>	<a href="#">6</a>
<a href="#">4.</a>	<a href="#">The DTLS Record Layer</a>	<a href="#">7</a>
<a href="#">4.1.</a>	<a href="#">Transport Layer Mapping</a>	<a href="#">8</a>
<a href="#">4.2.</a>	<a href="#">PMTU Issues</a>	<a href="#">9</a>
<a href="#">4.3.</a>	<a href="#">Record Payload Protection</a>	<a href="#">11</a>
<a href="#">4.3.1.</a>	<a href="#">Anti-Replay</a>	<a href="#">11</a>
<a href="#">4.3.2.</a>	<a href="#">Handling Invalid Records</a>	<a href="#">11</a>
<a href="#">5.</a>	<a href="#">The DTLS Handshake Protocol</a>	<a href="#">12</a>
<a href="#">5.1.</a>	<a href="#">Denial-of-Service Countermeasures</a>	<a href="#">12</a>
<a href="#">5.2.</a>	<a href="#">DTLS Handshake Message Format</a>	<a href="#">15</a>
<a href="#">5.3.</a>	<a href="#">ACK Message</a>	<a href="#">19</a>
<a href="#">5.4.</a>	<a href="#">Handshake Message Fragmentation and Reassembly</a>	<a href="#">19</a>
<a href="#">5.5.</a>	<a href="#">Timeout and Retransmission</a>	<a href="#">20</a>
<a href="#">5.5.1.</a>	<a href="#">State Machine</a>	<a href="#">24</a>
<a href="#">5.5.2.</a>	<a href="#">Timer Values</a>	<a href="#">27</a>
<a href="#">5.6.</a>	<a href="#">CertificateVerify and Finished Messages</a>	<a href="#">28</a>
<a href="#">5.7.</a>	<a href="#">Alert Messages</a>	<a href="#">28</a>
5.8.	<a href="#">Establishing New Associations with Existing Parameters</a>	28
<a href="#">5.9.</a>	<a href="#">Epoch Values and Rekeying</a>	<a href="#">28</a>
<a href="#">6.</a>	<a href="#">Application Data Protocol</a>	<a href="#">30</a>
<a href="#">7.</a>	<a href="#">Security Considerations</a>	<a href="#">30</a>
<a href="#">8.</a>	<a href="#">Changes to DTLS 1.2</a>	<a href="#">30</a>



<a href="#">9.</a>	<a href="#">Open Issues</a>	<a href="#">31</a>
<a href="#">10.</a>	<a href="#">IANA Considerations</a>	<a href="#">31</a>
<a href="#">11.</a>	<a href="#">References</a>	<a href="#">31</a>
<a href="#">11.1.</a>	<a href="#">Normative References</a>	<a href="#">31</a>
<a href="#">11.2.</a>	<a href="#">Informative References</a>	<a href="#">32</a>
<a href="#">Appendix A.</a>	<a href="#">History</a>	<a href="#">34</a>
<a href="#">Appendix B.</a>	<a href="#">Working Group Information</a>	<a href="#">34</a>
<a href="#">Appendix C.</a>	<a href="#">Contributors</a>	<a href="#">34</a>
	<a href="#">Authors' Addresses</a>	<a href="#">35</a>

## **[1.](#) Introduction**

RFC EDITOR: PLEASE REMOVE THE FOLLOWING PARAGRAPH

The source for this draft is maintained in GitHub. Suggested changes should be submitted as pull requests at <https://github.com/tlswg/dtls13-spec>. Instructions are on that page as well. Editorial changes can be managed in GitHub, but any substantive change should be discussed on the TLS mailing list.

The primary goal of the TLS protocol is to provide privacy and data integrity between two communicating peers. The TLS protocol is composed of two layers: the TLS Record Protocol and the TLS Handshake Protocol. However, TLS must run over a reliable transport channel - typically TCP [[RFC0793](#)].

There are applications that utilize UDP as a transport and to offer communication security protection for those applications the Datagram Transport Layer Security (DTLS) protocol has been designed. DTLS is deliberately designed to be as similar to TLS as possible, both to minimize new security invention and to maximize the amount of code and infrastructure reuse.

DTLS 1.0 was originally defined as a delta from TLS 1.1 and DTLS 1.2 was defined as a series of deltas to TLS 1.2. There is no DTLS 1.1; that version number was skipped in order to harmonize version numbers with TLS. This specification describes the most current version of the DTLS protocol aligning with the efforts around TLS 1.3.

Implementations that speak both DTLS 1.2 and DTLS 1.3 can interoperate with those that speak only DTLS 1.2 (using DTLS 1.2 of course), just as TLS 1.3 implementations can interoperate with TLS 1.2 (see [Appendix C](#) of [[I-D.ietf-tls-tls13](#)] for details). While backwards compatibility with DTLS 1.0 is possible the use of DTLS 1.0 is not recommended as explained in [Section 3.1.2 of RFC 7525](#) [[RFC7525](#)].



## 2. Conventions and Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [RFC2119].

The following terms are used:

- client: The endpoint initiating the TLS connection.
- connection: A transport-layer connection between two endpoints.
- endpoint: Either the client or server of the connection.
- handshake: An initial negotiation between client and server that establishes the parameters of their transactions.
- peer: An endpoint. When discussing a particular endpoint, "peer" refers to the endpoint that is remote to the primary subject of discussion.
- receiver: An endpoint that is receiving records.
- sender: An endpoint that is transmitting records.
- session: An association between a client and a server resulting from a handshake.
- server: The endpoint which did not initiate the TLS connection.

The reader is assumed to be familiar with the TLS 1.3 specification since this document defined as a delta from TLS 1.3.

## 3. DTLS Design Rational and Overview

The basic design philosophy of DTLS is to construct "TLS over datagram transport". Datagram transport does not require or provide reliable or in-order delivery of data. The DTLS protocol preserves this property for application data. Applications such as media streaming, Internet telephony, and online gaming use datagram transport for communication due to the delay-sensitive nature of transported data. The behavior of such applications is unchanged when the DTLS protocol is used to secure communication, since the DTLS protocol does not compensate for lost or re-ordered data traffic.

TLS cannot be used directly in datagram environments for the following five reasons:

1. TLS does not allow independent decryption of individual records. Because the integrity check indirectly depends on a sequence number, if record N is not received, then the integrity check on record N+1 will be based on the wrong sequence number and thus will fail. DTLS solves this problem by adding explicit sequence numbers.
2. The TLS handshake is a lock-step cryptographic handshake. Messages must be transmitted and received in a defined order; any other order is an error. Clearly, this is incompatible with reordering and message loss.
3. Not all TLS 1.3 handshake messages (such as the NewSessionTicket message) are acknowledged. Hence, a new acknowledgement message has to be added to detect message loss.
4. Handshake messages are potentially larger than any given datagram, thus creating the problem of IP fragmentation.
5. Datagram transport protocols, like UDP, are more vulnerable to denial of service attacks and require a return-routability check with the help of cookies to be integrated into the handshake. A detailed discussion of countermeasures can be found in [Section 5.1](#).

### **[3.1. Packet Loss](#)**

DTLS uses a simple retransmission timer to handle packet loss. Figure 1 demonstrates the basic concept, using the first phase of the DTLS handshake:




Figure 1: DTLS Retransmission Example.





Once the client has transmitted the ClientHello message, it expects to see a HelloRetryRequest from the server. However, if the server's message is lost, the client knows that either the ClientHello or the HelloRetryRequest has been lost and retransmits. When the server receives the retransmission, it knows to retransmit.

The server also maintains a retransmission timer and retransmits when that timer expires.

Note that timeout and retransmission do not apply to the HelloRetryRequest since this would require creating state on the server. The HelloRetryRequest is designed to be small enough that it will not itself be fragmented, thus avoiding concerns about interleaving multiple HelloRetryRequests.

#### **3.1.1. Reordering**

In DTLS, each handshake message is assigned a specific sequence number within that handshake. When a peer receives a handshake message, it can quickly determine whether that message is the next message it expects. If it is, then it processes it. If not, it queues it for future handling once all previous messages have been received.

#### **3.1.2. Message Size**

TLS and DTLS handshake messages can be quite large (in theory up to  $2^{24}-1$  bytes, in practice many kilobytes). By contrast, UDP datagrams are often limited to less than 1500 bytes if IP fragmentation is not desired. In order to compensate for this limitation, each DTLS handshake message may be fragmented over several DTLS records, each of which is intended to fit in a single IP datagram. Each DTLS handshake message contains both a fragment offset and a fragment length. Thus, a recipient in possession of all bytes of a handshake message can reassemble the original unfragmented message.

#### **3.2. Replay Detection**

DTLS optionally supports record replay detection. The technique used is the same as in IPsec AH/ESP, by maintaining a bitmap window of received records. Records that are too old to fit in the window and records that have previously been received are silently discarded. The replay detection feature is optional, since packet duplication is not always malicious, but can also occur due to routing errors. Applications may conceivably detect duplicate packets and accordingly modify their data transmission strategy.



#### 4. The DTLS Record Layer

The DTLS record layer is extremely similar to that of TLS 1.3. The only change is the inclusion of an explicit epoch and sequence number in the record. This sequence number allows the recipient to correctly verify the TLS MAC. The DTLS record format is shown below:

```
struct {  
    ContentType type;  
    ProtocolVersion version = { 254, 253 };  
    uint16 epoch;                // DTLS-related field  
    uint48 sequence_number;      // DTLS-related field  
    uint16 length;  
    opaque fragment[DTLSPlaintext.length];  
} DTLSPlaintext;
```

type: Identical to the type field in a TLS 1.3 record.

version: This specification re-uses the DTLS version 1.2 version number, namely { 254, 253 }. This field is deprecated and MUST be ignored for all purposes.

epoch: A counter value that is incremented on every cipher state change.

sequence\_number: The sequence number for this record.

length: Identical to the length field in a TLS 1.3 record.

fragment: Identical to the fragment field in a TLS 1.3 record.

DTLS uses an explicit sequence number, rather than an implicit one, carried in the sequence\_number field of the record. Sequence numbers are maintained separately for each epoch, with each sequence\_number initially being 0 for each epoch. For instance, if a handshake message from epoch 0 is retransmitted, it might have a sequence number after a message from epoch 1, even if the message from epoch 1 was transmitted first. Note that some care needs to be taken during the handshake to ensure that retransmitted messages use the right epoch and keying material.

If several handshakes are performed in close succession, there might be multiple records on the wire with the same sequence number but from different cipher states. The epoch field allows recipients to distinguish such packets. The epoch number is initially zero and is incremented each time keying material changes and a sender aims to rekey. More details are provided in [Section 5.9](#). In order to ensure that any given sequence/epoch pair is unique, implementations MUST



NOT allow the same epoch value to be reused within two times the TCP maximum segment lifetime.

Note that because DTLS records may be reordered, a record from epoch 1 may be received after epoch 2 has begun. In general, implementations SHOULD discard packets from earlier epochs, but if packet loss causes noticeable problems they MAY choose to retain keying material from previous epochs for up to the default MSL specified for TCP [[RFC0793](#)] to allow for packet reordering. (Note that the intention here is that implementers use the current guidance from the IETF for MSL, not that they attempt to interrogate the MSL that the system TCP stack is using.) Until the handshake has completed, implementations MUST accept packets from the old epoch.

Conversely, it is possible for records that are protected by the newly negotiated context to be received prior to the completion of a handshake. For instance, the server may send its Finished message and then start transmitting data. Implementations MAY either buffer or discard such packets, though when DTLS is used over reliable transports (e.g., SCTP), they SHOULD be buffered and processed once the handshake completes. Note that TLS's restrictions on when packets may be sent still apply, and the receiver treats the packets as if they were sent in the right order. In particular, it is still impermissible to send data prior to completion of the first handshake.

As in TLS, implementations MUST either abandon an association or re-key using a KeyUpdate message prior to allowing the sequence number to wrap.

Implementations MUST NOT allow the epoch to wrap, but instead MUST establish a new association, terminating the old association.

#### **4.1. Transport Layer Mapping**

Each DTLS record MUST fit within a single datagram. In order to avoid IP fragmentation, clients of the DTLS record layer SHOULD attempt to size records so that they fit within any PMTU estimates obtained from the record layer.

Note that unlike IPsec, DTLS records do not contain any association identifiers. Applications must arrange to multiplex between associations. With UDP, this is presumably done with the host/port number.

Multiple DTLS records may be placed in a single datagram. They are simply encoded consecutively. The DTLS record framing is sufficient to determine the boundaries. Note, however, that the first byte of



the datagram payload must be the beginning of a record. Records may not span datagrams.

Some transports, such as DCCP [[RFC4340](#)] provide their own sequence numbers. When carried over those transports, both the DTLS and the transport sequence numbers will be present. Although this introduces a small amount of inefficiency, the transport layer and DTLS sequence numbers serve different purposes; therefore, for conceptual simplicity, it is superior to use both sequence numbers.

Some transports provide congestion control for traffic carried over them. If the congestion window is sufficiently narrow, DTLS handshake retransmissions may be held rather than transmitted immediately, potentially leading to timeouts and spurious retransmission. When DTLS is used over such transports, care should be taken not to overrun the likely congestion window. [[RFC5238](#)] defines a mapping of DTLS to DCCP that takes these issues into account.

#### **4.2. PMTU Issues**

In general, DTLS's philosophy is to leave PMTU discovery to the application. However, DTLS cannot completely ignore PMTU for three reasons:

- The DTLS record framing expands the datagram size, thus lowering the effective PMTU from the application's perspective.
- In some implementations, the application may not directly talk to the network, in which case the DTLS stack may absorb ICMP [[RFC1191](#)] "Datagram Too Big" indications or ICMPv6 [[RFC4443](#)] "Packet Too Big" indications.
- The DTLS handshake messages can exceed the PMTU.

In order to deal with the first two issues, the DTLS record layer SHOULD behave as described below.

If PMTU estimates are available from the underlying transport protocol, they should be made available to upper layer protocols. In particular:

- For DTLS over UDP, the upper layer protocol SHOULD be allowed to obtain the PMTU estimate maintained in the IP layer.
- For DTLS over DCCP, the upper layer protocol SHOULD be allowed to obtain the current estimate of the PMTU.





- For DTLS over TCP or SCTP, which automatically fragment and reassemble datagrams, there is no PMTU limitation. However, the upper layer protocol MUST NOT write any record that exceeds the maximum record size of  $2^{14}$  bytes.

The DTLS record layer SHOULD allow the upper layer protocol to discover the amount of record expansion expected by the DTLS processing.

If there is a transport protocol indication (either via ICMP or via a refusal to send the datagram as in [Section 14 of \[RFC4340\]](#)), then the DTLS record layer MUST inform the upper layer protocol of the error.

The DTLS record layer SHOULD NOT interfere with upper layer protocols performing PMTU discovery, whether via [\[RFC1191\]](#) or [\[RFC4821\]](#) mechanisms. In particular:

- Where allowed by the underlying transport protocol, the upper layer protocol SHOULD be allowed to set the state of the DF bit (in IPv4) or prohibit local fragmentation (in IPv6).
- If the underlying transport protocol allows the application to request PMTU probing (e.g., DCCP), the DTLS record layer should honor this request.

The final issue is the DTLS handshake protocol. From the perspective of the DTLS record layer, this is merely another upper layer protocol. However, DTLS handshakes occur infrequently and involve only a few round trips; therefore, the handshake protocol PMTU handling places a premium on rapid completion over accurate PMTU discovery. In order to allow connections under these circumstances, DTLS implementations SHOULD follow the following rules:

- If the DTLS record layer informs the DTLS handshake layer that a message is too big, it SHOULD immediately attempt to fragment it, using any existing information about the PMTU.
- If repeated retransmissions do not result in a response, and the PMTU is unknown, subsequent retransmissions SHOULD back off to a smaller record size, fragmenting the handshake message as appropriate. This standard does not specify an exact number of retransmits to attempt before backing off, but 2-3 seems appropriate.



### **4.3. Record Payload Protection**

Like TLS, DTLS transmits data as a series of protected records. The rest of this section describes the details of that format.

#### **4.3.1. Anti-Replay**

DTLS records contain a sequence number to provide replay protection. Sequence number verification SHOULD be performed using the following sliding window procedure, borrowed from [Section 3.4.3 of \[RFC4303\]](#).

The receiver packet counter for this session MUST be initialized to zero when the session is established. For each received record, the receiver MUST verify that the record contains a sequence number that does not duplicate the sequence number of any other record received during the life of this session. This SHOULD be the first check applied to a packet after it has been matched to a session, to speed rejection of duplicate records.

Duplicates are rejected through the use of a sliding receive window. (How the window is implemented is a local matter, but the following text describes the functionality that the implementation must exhibit.) A minimum window size of 32 MUST be supported, but a window size of 64 is preferred and SHOULD be employed as the default. Another window size (larger than the minimum) MAY be chosen by the receiver. (The receiver does not notify the sender of the window size.)

The "right" edge of the window represents the highest validated sequence number value received on this session. Records that contain sequence numbers lower than the "left" edge of the window are rejected. Packets falling within the window are checked against a list of received packets within the window. An efficient means for performing this check, based on the use of a bit mask, is described in [Section 3.4.3 of \[RFC4303\]](#).

If the received record falls within the window and is new, or if the packet is to the right of the window, then the receiver proceeds to MAC verification. If the MAC validation fails, the receiver MUST discard the received record as invalid. The receive window is updated only if the MAC verification succeeds.

#### **4.3.2. Handling Invalid Records**

Unlike TLS, DTLS is resilient in the face of invalid records (e.g., invalid formatting, length, MAC, etc.). In general, invalid records SHOULD be silently discarded, thus preserving the association; however, an error MAY be logged for diagnostic purposes.



Implementations which choose to generate an alert instead, **MUST** generate fatal level alerts to avoid attacks where the attacker repeatedly probes the implementation to see how it responds to various types of error. Note that if DTLS is run over UDP, then any implementation which does this will be extremely susceptible to denial-of-service (DoS) attacks because UDP forgery is so easy. Thus, this practice is **NOT RECOMMENDED** for such transports.

If DTLS is being carried over a transport that is resistant to forgery (e.g., SCTP with SCTP-AUTH), then it is safer to send alerts because an attacker will have difficulty forging a datagram that will not be rejected by the transport layer.

## **5. The DTLS Handshake Protocol**

DTLS 1.3 re-uses the TLS 1.3 handshake messages and flows, with the following changes:

1. To handle message loss, reordering, and fragmentation modifications to the handshake header are necessary.
2. Retransmission timers are introduced to handle message loss.
3. The TLS 1.3 KeyUpdate message is not used in DTLS 1.3 for re-keying.
4. A new ACK message is introduced to more robustness in message delivery.

Note that TLS 1.3 already supports a cookie extension, which used to prevent denial-of-service attacks. This DoS prevention mechanism is described in more detail below since UDP-based protocols are more vulnerable to amplification attacks than a connection-oriented transport like TCP that performs return-routability checks as part of the connection establishment.

With these exceptions, the DTLS message formats, flows, and logic are the same as those of TLS 1.3.

### **5.1. Denial-of-Service Countermeasures**

Datagram security protocols are extremely susceptible to a variety of DoS attacks. Two attacks are of particular concern:

1. An attacker can consume excessive resources on the server by transmitting a series of handshake initiation requests, causing the server to allocate state and potentially to perform expensive cryptographic operations.



2. An attacker can use the server as an amplifier by sending connection initiation messages with a forged source of the victim. The server then sends its next message (in DTLS, a Certificate message, which can be quite large) to the victim machine, thus flooding it.

In order to counter both of these attacks, DTLS borrows the stateless cookie technique used by Photuris [[RFC2522](#)] and IKE [[RFC5996](#)]. When the client sends its ClientHello message to the server, the server MAY respond with a HelloRetryRequest message. The HelloRetryRequest message as well as the cookie extension is defined in TLS 1.3. The HelloRetryRequest message contains a stateless cookie generated using the technique of [[RFC2522](#)]. The client MUST retransmit the ClientHello with the cookie added as an extension. The server then verifies the cookie and proceeds with the handshake only if it is valid. This mechanism forces the attacker/client to be able to receive the cookie, which makes DoS attacks with spoofed IP addresses difficult. This mechanism does not provide any defence against DoS attacks mounted from valid IP addresses.

The DTLS 1.3 specification changes the way how cookies are exchanged compared to DTLS 1.2. DTLS 1.3 re-uses the HelloRetryRequest message and conveys the cookie to the client via an extension. The client then uses the same extension to place the cookie into a ClientHello message. DTLS 1.2 on the other hand used a separate message, namely the HelloVerifyRequest, to pass a cookie to the client and did not utilize the extension mechanism. For backwards compatibility reason the cookie field in the ClientHello is present in DTLS 1.3 but is ignored by a DTLS 1.3 compliant server implementation.

The exchange is shown in Figure 2. Note that the figure focuses on the cookie exchange; all other extensions are omitted.

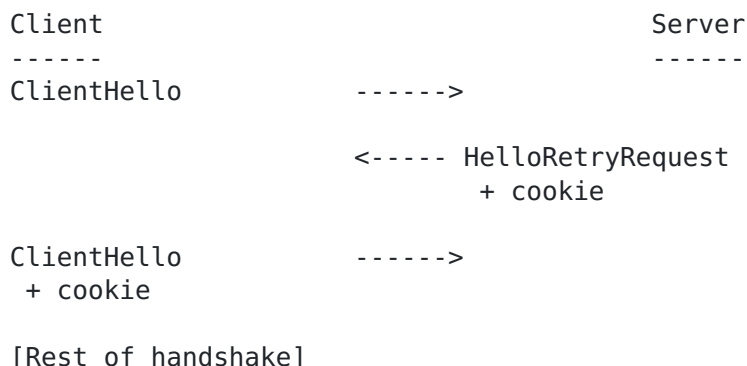


Figure 2: DTLS Exchange with HelloRetryRequest contain the Cookie Extension





The cookie extension is defined in Section 4.2.1 of [\[I-D.ietf-tls-tls13\]](#). When sending the initial ClientHello, the client does not have a cookie yet. In this case, the cookie extension is omitted and the legacy\_cookie field in the ClientHello message SHOULD be set to a zero length vector (i.e., a single zero byte length field) and MUST be ignored by a server negotiating DTLS 1.3.

When responding to a HelloRetryRequest, the client MUST create a new ClientHello message following the description in Section 4.1 of [\[I-D.ietf-tls-tls13\]](#).

The server SHOULD use information received in the ClientHello to generate its cookie, such as version, random, ciphersuites. The server MUST use the same version number in the HelloRetryRequest that it would use when sending a ServerHello. Upon receipt of the ServerHello, the client MUST verify that the server version values match.

If the HelloRetryRequest message is used, the initial ClientHello and the HelloRetryRequest are included in the calculation of the handshake\_messages (for the CertificateVerify message) and verify\_data (for the Finished message).

As such, the handshake transcript is not reset with the second ClientHello and a stateless server-cookie implementation requires the transcript of the HelloRetryRequest to be stored in the cookie or the internal state of the hash algorithm, since only the hash of the transcript is required for the handshake to complete.

When the second ClientHello is received, the server can verify that the cookie is valid and that the client can receive packets at the given IP address.

One potential attack on this scheme is for the attacker to collect a number of cookies from different addresses and then reuse them to attack the server. The server can defend against this attack by changing the secret value frequently, thus invalidating those cookies. If the server wishes that legitimate clients be able to handshake through the transition (e.g., they received a cookie with Secret 1 and then sent the second ClientHello after the server has changed to Secret 2), the server can have a limited window during which it accepts both secrets. [\[RFC5996\]](#) suggests adding a key identifier to cookies to detect this case. An alternative approach is simply to try verifying with both secrets. It is RECOMMENDED that servers implement a key rotation scheme that allows the server to manage keys with overlapping lifetime.



Alternatively, the server can store timestamps in the cookie and reject those cookies that were not generated within a certain amount of time.

DTLS servers SHOULD perform a cookie exchange whenever a new handshake is being performed. If the server is being operated in an environment where amplification is not a problem, the server MAY be configured not to perform a cookie exchange. The default SHOULD be that the exchange is performed, however. In addition, the server MAY choose not to do a cookie exchange when a session is resumed. Clients MUST be prepared to do a cookie exchange with every handshake.

If a server receives a ClientHello with an invalid cookie, it MUST NOT respond with a HelloRetryRequest. Restarting the handshake from scratch, without a cookie, allows the client to recover from a situation where it obtained a cookie that cannot be verified by the server. As described in Section 4.1.4 of [\[I-D.ietf-tls-tls13\]](#), clients SHOULD also abort the handshake with an "unexpected\_message" alert in response to any second HelloRetryRequest which was sent in the same connection (i.e., where the ClientHello was itself in response to a HelloRetryRequest).

## **5.2. DTLS Handshake Message Format**

In order to support message loss, reordering, and message fragmentation, DTLS modifies the TLS 1.3 handshake header:

```
enum {
    hello_request_RESERVED(0),
    client_hello(1),
    server_hello(2),
    hello_verify_request_RESERVED(3),
    new_session_ticket(4),
    hello_retry_request(6),
    encrypted_extensions(8),
    certificate(11),
    server_key_exchange_RESERVED(12),
    certificate_request(13),
    server_hello_done_RESERVED(14),
    certificate_verify(15),
    client_key_exchange_RESERVED(16),
    finished(20),
    key_update_RESERVED(24),
    (255)
} HandshakeType;

struct {
    HandshakeType msg_type;      /* handshake type */
    uint24 length;              /* bytes in message */
    uint16 message_seq;         /* DTLS-required field */
    uint24 fragment_offset;     /* DTLS-required field */
    uint24 fragment_length;     /* DTLS-required field */
    select (HandshakeType) {
        case client_hello:      ClientHello;
        case server_hello:      ServerHello;
        case hello_retry_request: HelloRetryRequest;
        case encrypted_extensions: EncryptedExtensions;
        case certificate_request: CertificateRequest;
        case certificate:        Certificate;
        case certificate_verify: CertificateVerify;
        case finished:           Finished;
        case new_session_ticket: NewSessionTicket;
        case key_update:         KeyUpdate; /* reserved */
        case ack:                ACK; /* DTLS-required field */
    } body;
} Handshake;
```

In addition to the handshake messages that are deprecated by the TLS 1.3 specification DTLS 1.3 furthermore deprecates the HelloVerifyRequest message originally defined in DTLS 1.0. DTLS 1.3-compliant implements MUST NOT use the HelloVerifyRequest to execute a return-routability check. A dual-stack DTLS 1.2/DTLS 1.3 client must, however, be prepared to interact with a DTLS 1.2 server.



A DTLS 1.3 MUST NOT use the KeyUpdate message to change keying material used for the protection of traffic data. Instead the epoch field is used, which is explained in [Section 5.9](#).

The format of the ClientHello used by a DTLS 1.3 client differs from the TLS 1.3 ClientHello format as shown below.

```
struct {  
    ProtocolVersion client_version = { 254,252 };    /* DTLS v1.3 */  
    Random random;  
    opaque legacy_session_id<0..32>;  
    opaque legacy_cookie<0..2^8-1>;                // DTLS  
    CipherSuite cipher_suites<2..2^16-2>;  
    opaque legacy_compression_methods<1..2^8-1>;  
    Extension extensions<0..2^16-1>;  
} ClientHello;
```

**client\_version:** The version of the DTLS protocol by which the client wishes to communicate during this session. This SHOULD be the latest (highest valued) version supported by the client. For the DTLS 1.3 version of the specification, the version will be { 254,252 }.

**random:** Same as for TLS 1.3

**legacy\_session\_id:** Same as for TLS 1.3

**legacy\_cookie:** A DTLS 1.3-only client MUST set the legacy\_cookie field to zero length.

**cipher\_suites:** Same as for TLS 1.3

**legacy\_compression\_methods:** Same as for TLS 1.3

**extensions:** Same as for TLS 1.3

The first message each side transmits in each handshake always has message\_seq = 0. Whenever a new message is generated, the message\_seq value is incremented by one. When a message is retransmitted, the old message\_seq value is re-used, i.e., not incremented.

Here is an example:



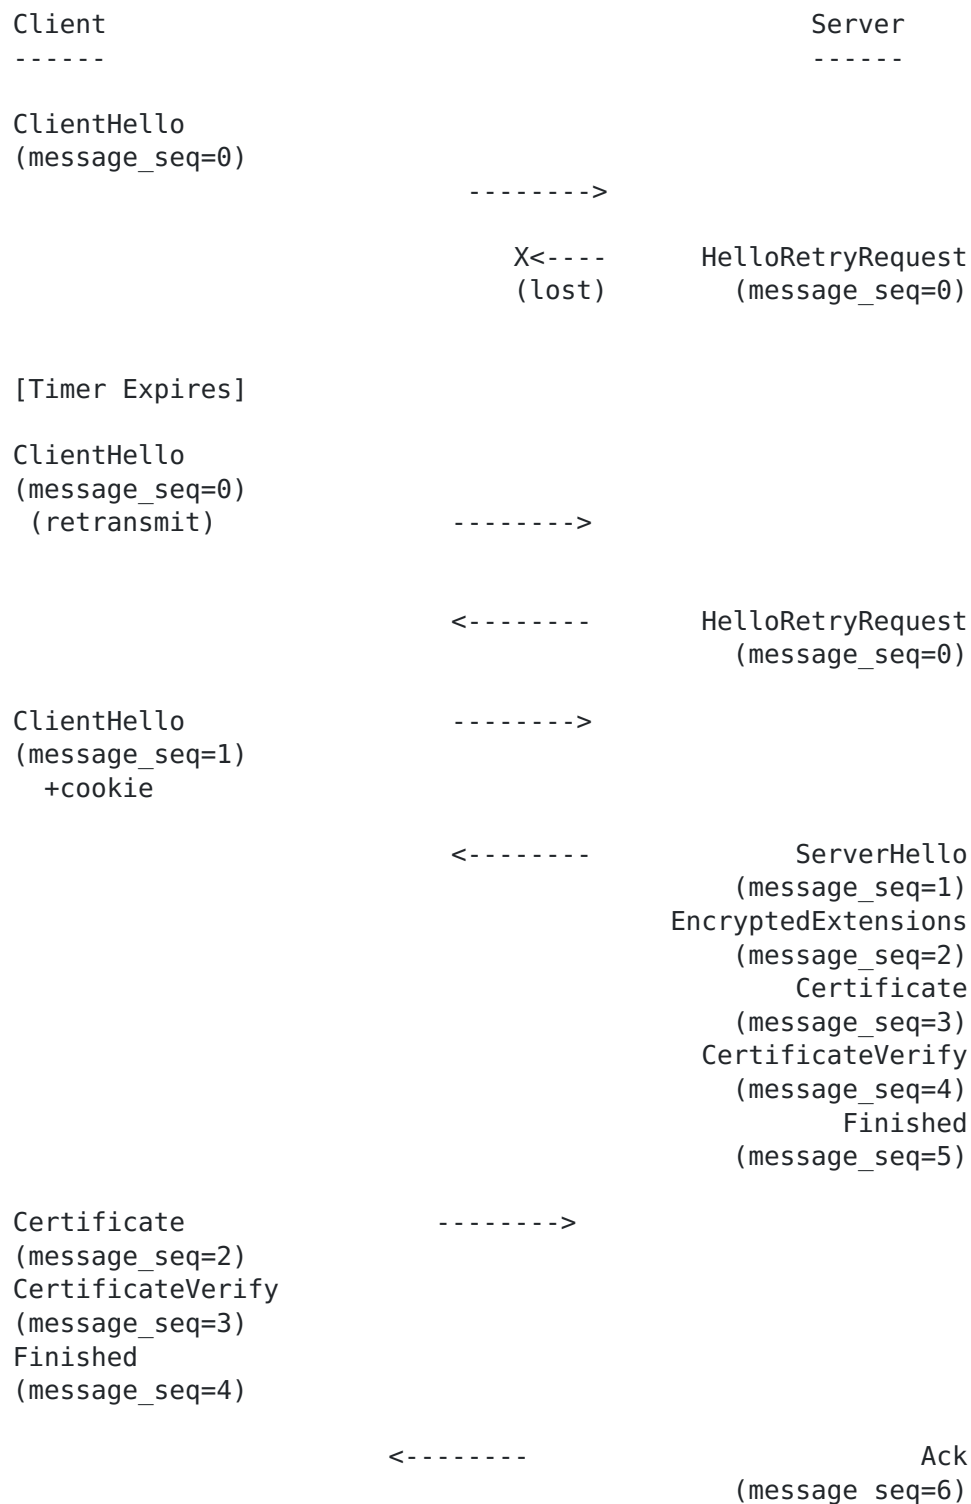


Figure 3: Example DTLS Exchange illustrating Message Loss





From the perspective of the DTLS record layer, the retransmission is a new record. This record will have a new DTLSPlaintext.sequence\_number value.

DTLS implementations maintain (at least notionally) a next\_receive\_seq counter. This counter is initially set to zero. When a message is received, if its sequence number matches next\_receive\_seq, next\_receive\_seq is incremented and the message is processed. If the sequence number is less than next\_receive\_seq, the message MUST be discarded. If the sequence number is greater than next\_receive\_seq, the implementation SHOULD queue the message but MAY discard it. (This is a simple space/bandwidth tradeoff).

### **5.3. ACK Message**

```
struct {} ACK;
```

The ACK handshake message is used by a server to return a response to a client-provided message where the TLS 1.3 handshake does not foresee such return message. With the use of the ACK message the client is able to determine whether a transmitted request has been lost and needs to be retransmitted. Since the ACK message does not contain any correlation information the server MUST only have one message outstanding at a time.

### **5.4. Handshake Message Fragmentation and Reassembly**

Each DTLS message MUST fit within a single transport layer datagram. However, handshake messages are potentially bigger than the maximum record size. Therefore, DTLS provides a mechanism for fragmenting a handshake message over a number of records, each of which can be transmitted separately, thus avoiding IP fragmentation.

When transmitting the handshake message, the sender divides the message into a series of N contiguous data ranges. These ranges MUST NOT be larger than the maximum handshake fragment size and MUST jointly contain the entire handshake message. The ranges MUST NOT overlap. The sender then creates N handshake messages, all with the same message\_seq value as the original handshake message. Each new message is labeled with the fragment\_offset (the number of bytes contained in previous fragments) and the fragment\_length (the length of this fragment). The length field in all messages is the same as the length field of the original message. An unfragmented message is a degenerate case with fragment\_offset=0 and fragment\_length=length.

When a DTLS implementation receives a handshake message fragment, it MUST buffer it until it has the entire handshake message. DTLS implementations MUST be able to handle overlapping fragment ranges.



This allows senders to retransmit handshake messages with smaller fragment sizes if the PMTU estimate changes.

Note that as with TLS, multiple handshake messages may be placed in the same DTLS record, provided that there is room and that they are part of the same flight. Thus, there are two acceptable ways to pack two DTLS messages into the same datagram: in the same record or in separate records.

### **5.5. Timeout and Retransmission**

DTLS messages are grouped into a series of message flights, according to the diagrams below. Although each flight of messages may consist of a number of messages, they should be viewed as monolithic for the purpose of timeout and retransmission.

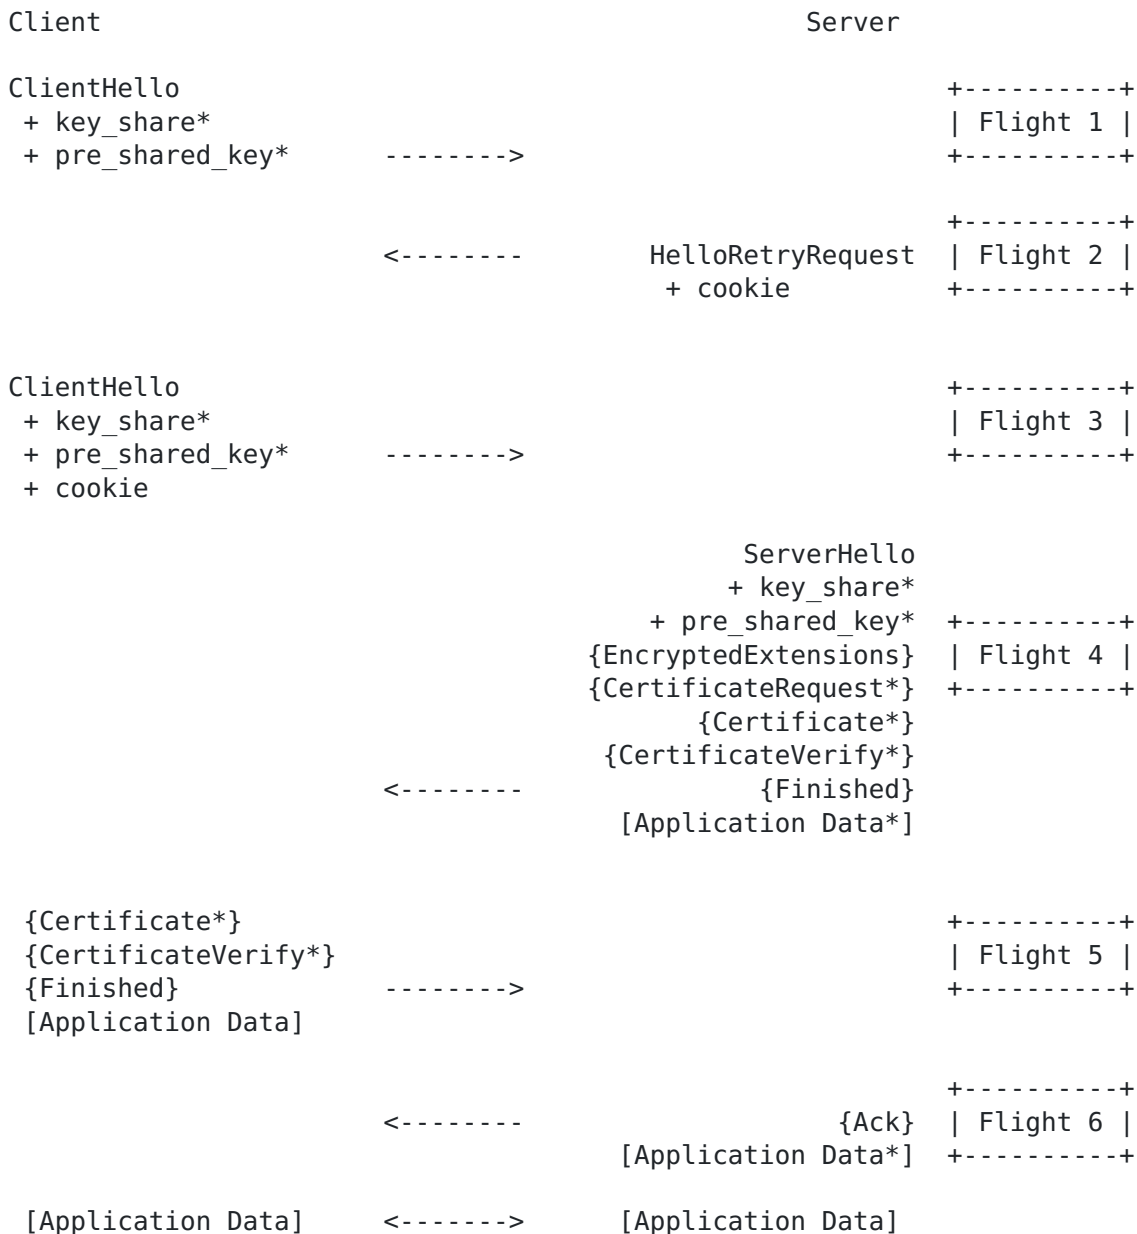


Figure 4: Message Flights for full DTLS Handshake (with Cookie Exchange)

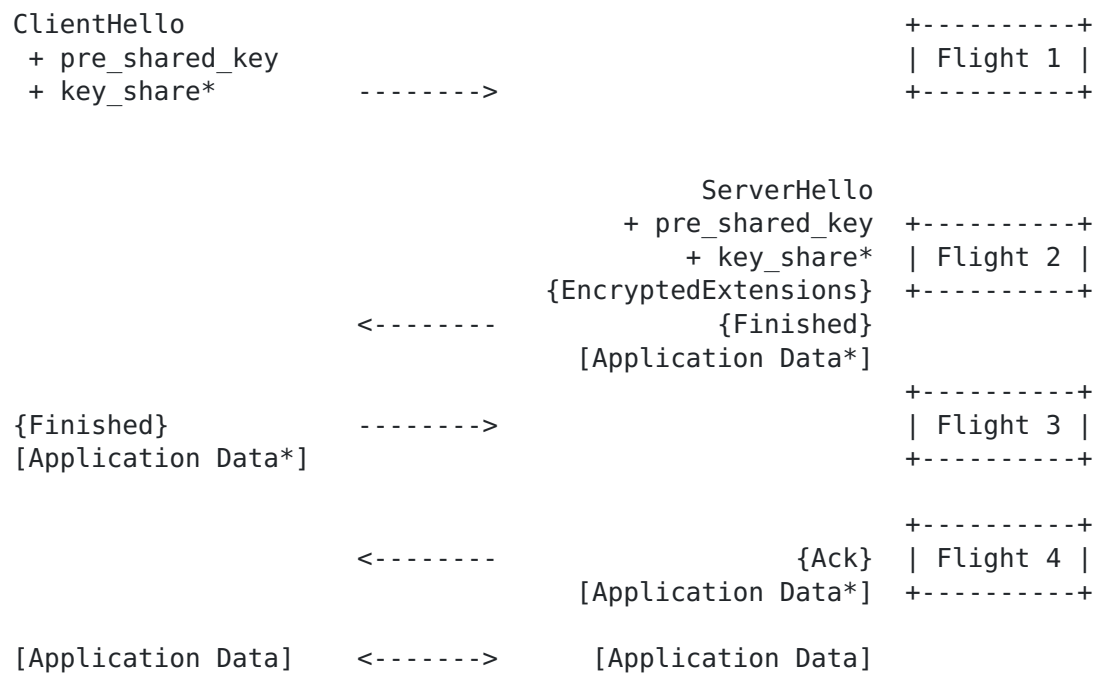


Figure 5: Message Flights for Resumption and PSK Handshake (without Cookie Exchange)

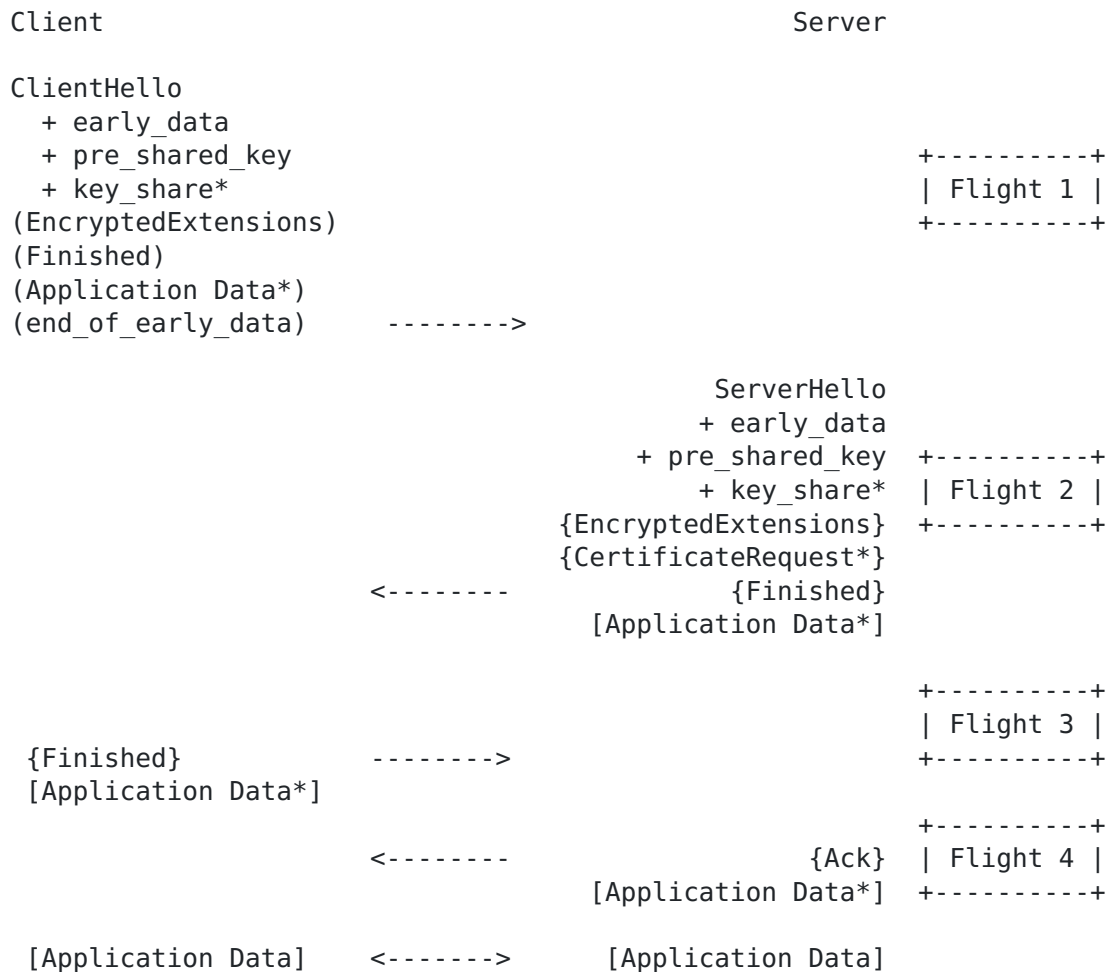


Figure 6: Message Flights for a zero round trip handshake

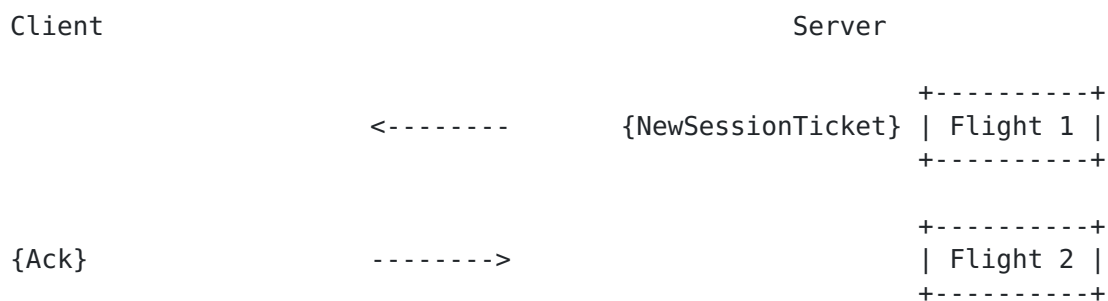


Figure 7: Message Flights for New Session Ticket Message








Figure 8: Message Flights for Post-Handshake Authentication (Success)

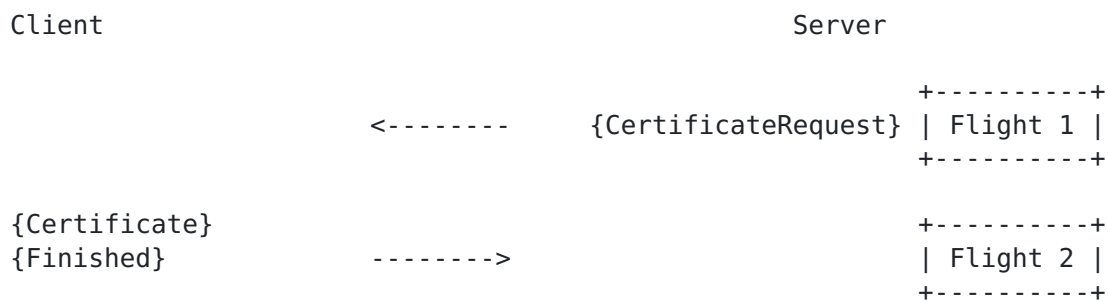


Figure 9: Message Flights for Post-Handshake Authentication (Decline)

Note: The application data sent by the client is not included in the timeout and retransmission calculation.

#### 5.5.1. State Machine

DTLS uses a simple timeout and retransmission scheme with the state machine shown in Figure 10. Because DTLS clients send the first message (`ClientHello`), they start in the `PREPARING` state. DTLS servers start in the `WAITING` state, but with empty buffers and no retransmit timer.

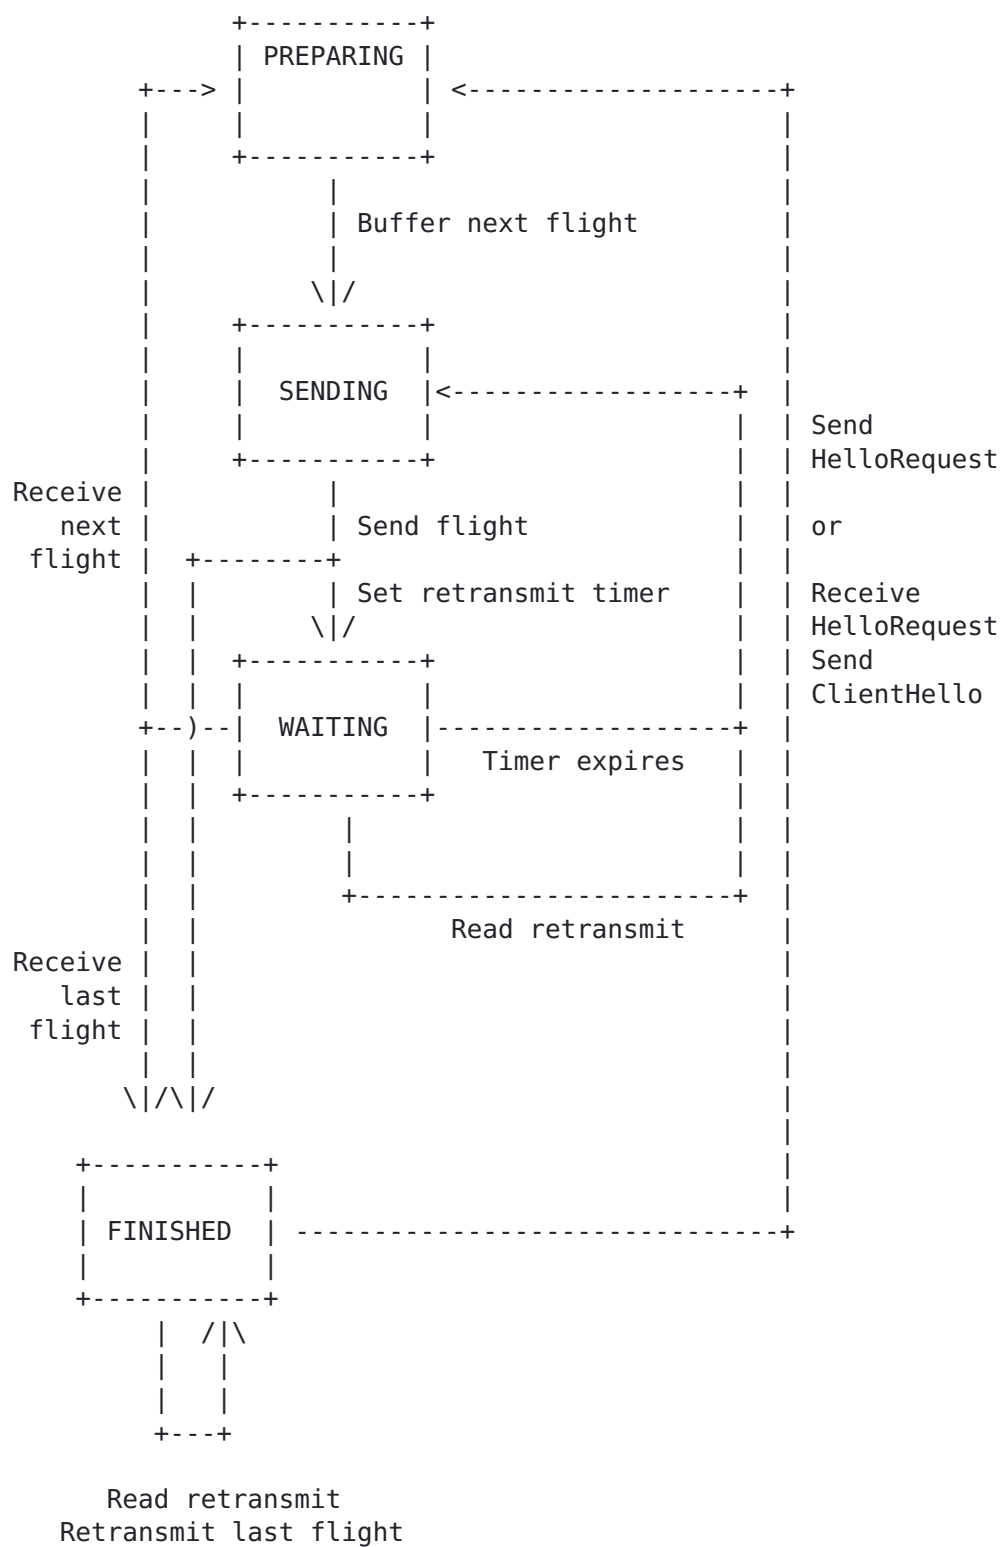


Figure 10: DTLS Timeout and Retransmission State Machine



The state machine has three basic states.

In the PREPARING state, the implementation does whatever computations are necessary to prepare the next flight of messages. It then buffers them up for transmission (emptying the buffer first) and enters the SENDING state.

In the SENDING state, the implementation transmits the buffered flight of messages. Once the messages have been sent, the implementation then enters the FINISHED state if this is the last flight in the handshake. Or, if the implementation expects to receive more messages, it sets a retransmit timer and then enters the WAITING state.

There are three ways to exit the WAITING state:

1. The retransmit timer expires: the implementation transitions to the SENDING state, where it retransmits the flight, resets the retransmit timer, and returns to the WAITING state.
2. The implementation reads a retransmitted flight from the peer: the implementation transitions to the SENDING state, where it retransmits the flight, resets the retransmit timer, and returns to the WAITING state. The rationale here is that the receipt of a duplicate message is the likely result of timer expiry on the peer and therefore suggests that part of one's previous flight was lost.
3. The implementation receives the next flight of messages: if this is the final flight of messages, the implementation transitions to FINISHED. If the implementation needs to send a new flight, it transitions to the PREPARING state. Partial reads (whether partial messages or only some of the messages in the flight) do not cause state transitions or timer resets.

Because DTLS clients send the first message (ClientHello), they start in the PREPARING state. DTLS servers start in the WAITING state, but with empty buffers and no retransmit timer.

When the server desires a rehandshake, it transitions from the FINISHED state to the PREPARING state to transmit the HelloRequest. When the client receives a HelloRequest, it transitions from FINISHED to PREPARING to transmit the ClientHello.

In addition, for at least twice the default Maximum Segment Lifetime (MSL) defined for [RFC0793](#), when in the FINISHED state, the node that transmits the last flight (the server in an



ordinary handshake or the client in a resumed handshake) MUST respond to a retransmit of the peer's last flight with a retransmit of the last flight. This avoids deadlock conditions if the last flight gets lost. To see why this is necessary, consider what happens in an ordinary handshake if the server's Finished message is lost: the server believes the handshake is complete but it actually is not. As the client is waiting for the Finished message, the client's retransmit timer will fire and it will retransmit the client's Finished message. This will cause the server to respond with its own Finished message, completing the handshake. The same logic applies on the server side for the resumed handshake.

Note that because of packet loss, it is possible for one side to be sending application data even though the other side has not received the first side's Finished message. Implementations MUST either discard or buffer all application data packets for the new epoch until they have received the Finished message for that epoch. Implementations MAY treat receipt of application data with a new epoch prior to receipt of the corresponding Finished message as evidence of reordering or packet loss and retransmit their final flight immediately, shortcutting the retransmission timer.

#### **5.5.2. Timer Values**

Though timer values are the choice of the implementation, mishandling of the timer can lead to serious congestion problems; for example, if many instances of a DTLS time out early and retransmit too quickly on a congested link. Implementations SHOULD use an initial timer value of 100 msec (the minimum defined in [RFC 6298](#) [RFC6298]) and double the value at each retransmission, up to no less than the [RFC 6298](#) maximum of 60 seconds. Application specific profiles, such as those used for the Internet of Things environment, may recommend longer timer values. Note that we recommend a 100 msec timer rather than the 3-second [RFC 6298](#) default in order to improve latency for time-sensitive applications. Because DTLS only uses retransmission for handshake and not dataflow, the effect on congestion should be minimal.

Implementations SHOULD retain the current timer value until a transmission without loss occurs, at which time the value may be reset to the initial value. After a long period of idleness, no less than 10 times the current timer value, implementations may reset the timer to the initial value. One situation where this might occur is when a rehandshake is used after substantial data transfer.



### **5.6. CertificateVerify and Finished Messages**

CertificateVerify and Finished messages have the same format as in TLS 1.3. Hash calculations include entire handshake messages, including DTLS-specific fields: message\_seq, fragment\_offset, and fragment\_length. However, in order to remove sensitivity to handshake message fragmentation, the CertificateVerify and the Finished messages MUST be computed as if each handshake message had been sent as a single fragment following the algorithm described in [Section 4.4.1](#) and Section 4.4.3 of [\[I-D.ietf-tls-tls13\]](#), respectively.

### **5.7. Alert Messages**

Note that Alert messages are not retransmitted at all, even when they occur in the context of a handshake. However, a DTLS implementation which would ordinarily issue an alert SHOULD generate a new alert message if the offending record is received again (e.g., as a retransmitted handshake message). Implementations SHOULD detect when a peer is persistently sending bad messages and terminate the local connection state after such misbehavior is detected.

### **5.8. Establishing New Associations with Existing Parameters**

If a DTLS client-server pair is configured in such a way that repeated connections happen on the same host/port quartet, then it is possible that a client will silently abandon one connection and then initiate another with the same parameters (e.g., after a reboot). This will appear to the server as a new handshake with epoch=0. In cases where a server believes it has an existing association on a given host/port quartet and it receives an epoch=0 ClientHello, it SHOULD proceed with a new handshake but MUST NOT destroy the existing association until the client has demonstrated reachability either by completing a cookie exchange or by completing a complete handshake including delivering a verifiable Finished message. After a correct Finished message is received, the server MUST abandon the previous association to avoid confusion between two valid associations with overlapping epochs. The reachability requirement prevents off-path/blind attackers from destroying associations merely by sending forged ClientHellos.

### **5.9. Epoch Values and Rekeying**

A recipient of a DTLS message needs to select the correct keying material in order to process an incoming message. With the possibility of message loss and re-order an identifier is needed to determine which cipher state has been used to protect the record payload. The epoch value fulfills this role in DTLS. In addition to





the key derivation steps described in Section 7 of [\[I-D.ietf-tls-tls13\]](#) triggered by the states during the handshake a sender may want to rekey at any time during the lifetime of the connection and has to have a way to indicate that it is updating its sending cryptographic keys.

The following epoch values are reserved and correspond to phases in the handshake and allow identification of the correct cipher state:

- epoch value (0) for use with unencrypted messages, namely ClientHello, ServerHello, and HelloRetryRequest.
- epoch value (1) for messages protected using keys derived from `early_traffic_secret`.
- epoch value (2) for 0-RTT 'Application Data' protected using keys derived from the `early_traffic_secret`.
- epoch value (3) for messages protected using keys derived from the `handshake_traffic_secret`, namely the EncryptedExtensions, CertificateRequest, Certificate, CertificateVerify, Finished, ACK, and NewSessionTicket messages).
- epoch value (4) for application data payloads protected using keys derived from the `initial_traffic_secret_0`.
- epoch value (5 to  $2^{16}-1$ ) for application data payloads protected using keys from the `traffic_secret_N` ( $N>0$ ).

Using these reserved epoch values a receiver knows what cipher state has been used to encrypt and integrity protect a message. Implementations that receive a payload with an epoch value for which no corresponding cipher state can be determined MUST generate a fatal "unexpected\_message" alert. For example, client incorrectly uses epoch value 5 when sending application data in a 0-RTT exchange with the first message. A server will not be able to compute the appropriate keys and will therefore have to respond with a fatal alert.

Increasing the epoch value by a sender (starting with value 5 upwards) corresponds semantically to rekeying using the KeyUpdate message in TLS 1.3. Instead of utilizing an dedicated message in DTLS 1.3 the sender uses an increase in the epoch value to signal rekeying. Hence, a sender that decides to increment the epoch value (with value starting at 5) MUST send all its traffic using the next generation of keys, computed as described in Section 7.2 of [\[I-D.ietf-tls-tls13\]](#). Upon receiving a payload with such a new epoch value, the receiver MUST update their receiving keys and if they have



not already updated their sending state up to or past the then current receiving generation MUST send messages with the new epoch value prior to sending any other messages. For epoch values lower than 5 the key schedule described in Section 7.1 of [\[I-D.ietf-tls-tls13\]](#) is applicable.

Note that epoch values do not wrap. If a DTLS implementation would need to wrap the epoch value, it MUST terminate the connection.

The traffic key calculation is described in Section 7.3 of [\[I-D.ietf-tls-tls13\]](#).

## **6. Application Data Protocol**

Application data messages are carried by the record layer and are fragmented and encrypted based on the current connection state. The messages are treated as transparent data to the record layer.

## **7. Security Considerations**

Security issues are discussed primarily in [\[I-D.ietf-tls-tls13\]](#).

The primary additional security consideration raised by DTLS is that of denial of service. DTLS includes a cookie exchange designed to protect against denial of service. However, implementations that do not use this cookie exchange are still vulnerable to DoS. In particular, DTLS servers that do not use the cookie exchange may be used as attack amplifiers even if they themselves are not experiencing DoS. Therefore, DTLS servers SHOULD use the cookie exchange unless there is good reason to believe that amplification is not a threat in their environment. Clients MUST be prepared to do a cookie exchange with every handshake.

Unlike TLS implementations, DTLS implementations SHOULD NOT respond to invalid records by terminating the connection.

## **8. Changes to DTLS 1.2**

Since TLS 1.3 introduce a large number of changes to TLS 1.2, the list of changes from DTLS 1.2 to DTLS 1.3 is equally large. For this reason this section focuses on the most important changes only.

- New handshake pattern, which leads to a shorter message exchange
- Support for AEAD-only ciphers
- HelloRetryRequest of TLS 1.3 used instead of HelloVerifyRequest



- More flexible ciphersuite negotiation
- New session resumption mechanism
- PSK authentication redefined
- New key derivation hierarchy utilizing the HKDF construct
- Removed support for weaker and older cryptographic algorithms

## **9. Open Issues**

- Handling of the handshake sequence numbers (i.e., `Handshake.message_seq`) when 0-RTT is rejected. Proposal: keep pushing the numbers forward
- Explore whether the record layer header can be simplified (to 2 octets for epoch & sequence number)
- Do we need the HelloRequest message in DTLS 1.3?
- Update text in the appendix regarding backwards compatibility.

## **10. IANA Considerations**

IANA is requested to allocate a new value in the TLS HandshakeType Registry for the ACK message defined in [Section 5.3](#).

## **11. References**

### **11.1. Normative References**

- [I-D.ietf-tls-tls13]  
Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", [draft-ietf-tls-tls13-18](#) (work in progress), October 2016.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), DOI 10.17487/RFC0793, September 1981, <<http://www.rfc-editor.org/info/rfc793>>.
- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", [RFC 1191](#), DOI 10.17487/RFC1191, November 1990, <<http://www.rfc-editor.org/info/rfc1191>>.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), DOI 10.17487/RFC4443, March 2006, <<http://www.rfc-editor.org/info/rfc4443>>.
- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", [RFC 4821](#), DOI 10.17487/RFC4821, March 2007, <<http://www.rfc-editor.org/info/rfc4821>>.
- [RFC6298] Paxson, V., Allman, M., Chu, J., and M. Sargent, "Computing TCP's Retransmission Timer", [RFC 6298](#), DOI 10.17487/RFC6298, June 2011, <<http://www.rfc-editor.org/info/rfc6298>>.

## **11.2. Informative References**

- [RFC2522] Karn, P. and W. Simpson, "Photuris: Session-Key Management Protocol", [RFC 2522](#), DOI 10.17487/RFC2522, March 1999, <<http://www.rfc-editor.org/info/rfc2522>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), DOI 10.17487/RFC4303, December 2005, <<http://www.rfc-editor.org/info/rfc4303>>.
- [RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", [RFC 4340](#), DOI 10.17487/RFC4340, March 2006, <<http://www.rfc-editor.org/info/rfc4340>>.
- [RFC5238] Phelan, T., "Datagram Transport Layer Security (DTLS) over the Datagram Congestion Control Protocol (DCCP)", [RFC 5238](#), DOI 10.17487/RFC5238, May 2008, <<http://www.rfc-editor.org/info/rfc5238>>.
- [RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", [RFC 5996](#), DOI 10.17487/RFC5996, September 2010, <<http://www.rfc-editor.org/info/rfc5996>>.





- [RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", [BCP 195](#), [RFC 7525](#), DOI 10.17487/RFC7525, May 2015, <<http://www.rfc-editor.org/info/rfc7525>>.

### **[11.3.](#) URIs**

- [1] <mailto:tls@ietf.org>

## **Appendix A. History**

RFC EDITOR: PLEASE REMOVE THE THIS SECTION

### **draft-00**

- Initial version using TLS 1.3 as a baseline.
- Use of epoch values instead of KeyUpdate message
- Use of cookie extension instead of cookie field in ClientHello and HelloVerifyRequest messages
- Added ACK message
- Text about sequence number handling

## **Appendix B. Working Group Information**

The discussion list for the IETF TLS working group is located at the e-mail address [tls@ietf.org](mailto:tls@ietf.org) [1]. Information on the group and information on how to subscribe to the list is at <https://www1.ietf.org/mailman/listinfo/tls>

Archives of the list can be found at: <https://www.ietf.org/mail-archive/web/tls/current/index.html>

## **Appendix C. Contributors**

Many people have contributed to previous DTLS versions and they are acknowledged in prior versions of DTLS specifications.

For this version of the document we would like to thank:

- \* Nagendra Modadugu (co-author of {{RFC6347}})  
Google, Inc.  
[nagendra@cs.stanford.edu](mailto:nagendra@cs.stanford.edu)
- \* Ilari Liusvaara  
Independent  
[ilariliusvaara@welho.com](mailto:ilariliusvaara@welho.com)
- \* Martin Thomson  
Mozilla  
[martin.thomson@gmail.com](mailto:martin.thomson@gmail.com)

Authors' Addresses

Eric Rescorla  
RTFM, Inc.

EMail: [ekr@rtfm.com](mailto:ekr@rtfm.com)

Hannes Tschofenig  
ARM Limited

EMail: [hannes.tschofenig@gmx.net](mailto:hannes.tschofenig@gmx.net)