DPRIVE Working Group
Internet-Draft
Intended status: Informational
Expires: September 16, 2016

D. Wing
T. Reddy
Cisco

March 15, 2016

DPRIVE TLS/DTLS Message Flows
draft-wing-dprive-profile-and-msg-flows-01

Abstract

Message flows for DNS-over-TLS and DNS-over-DTLS are discussed and compared.

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 September 16, 2016.

Table of Contents

1. Introduction

The DPRIVE working group has two active documents that provide DNS confidentiality, DNS over DTLS [I-D.ietf-dprive-dnsodtls] and DNS over TLS [I-D.ietf-dprive-dns-over-tls].

This document shows message flows for those two documents. Also shown is how TCP Fast Open (TFO) [RFC7413] eliminates a round-trip, which is especially helpful for DNS communication.

2. Server state lost

This section shows message flows after server state is lost, such as due to routing change (communicating to a different server, unbeknownst to the client) or due to server losing state (crash or software upgrade).

2.1. TLS

With TLS, the client is immediately informed of server state loss with a TCP RST, as shown in the diagram below. This costs one round trip, and this round trip is unavoidable. This is a TCP RST, and is not authenticated. After the RST, a new TCP connection and TLS state are established.

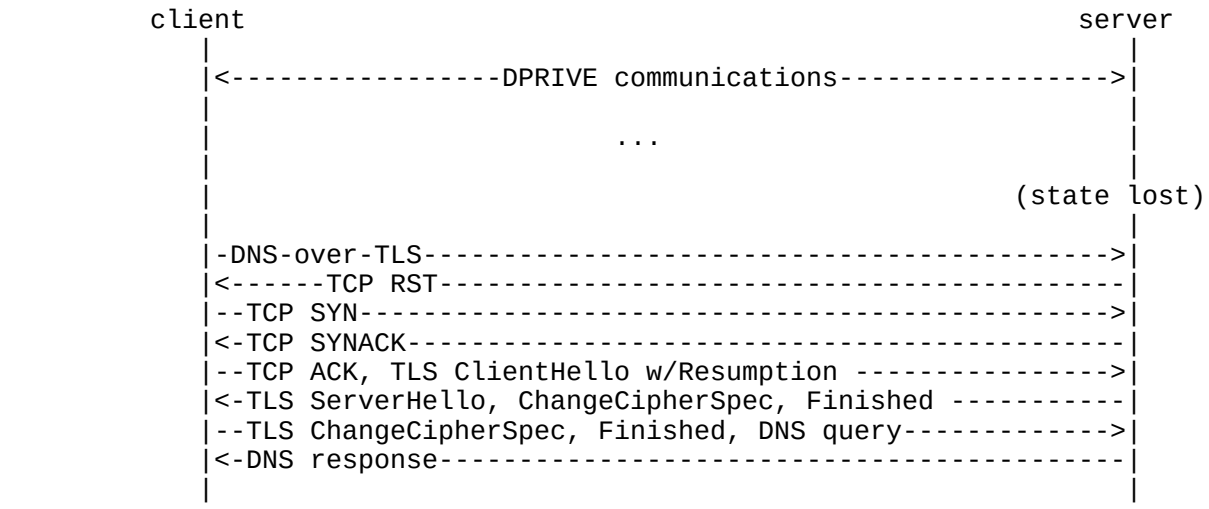

Figure 1: Server State Loss, TLS

2.2. DTLS

With DTLS, the client is immediately informed of the server state loss with a DTLS Alert, as shown in the diagram below. This DTLS Alert is not authenticated. This message costs one round trip, but can be avoided if the client anticipates this server state loss and consumes additional packet overhead, as discussed below Figure 2.

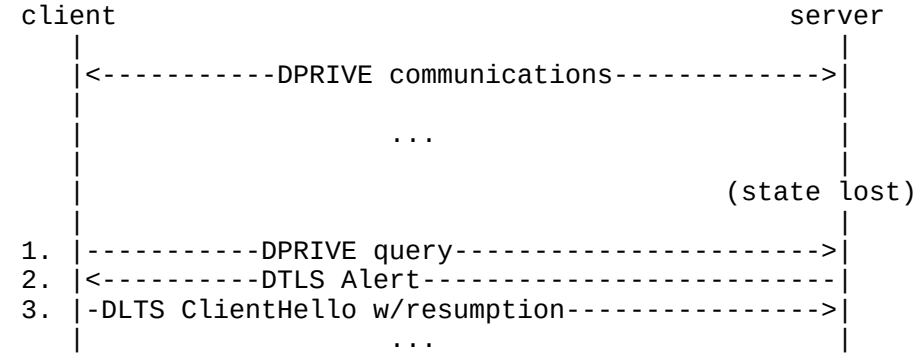

Figure 2: Server State Loss, DTLS

An optimization of the above flow is possible, if the client believes the server is likely to have lost state, such as if the most recent DPRIVE communications was a long time ago (exact value of "long time" is debatable). In that situation, the client can send a DTLS handshake with TLS resumption -- effectively, it sends the DTLS

handshake identical to packet (3) of Figure 2 (avoiding packets 1 and 2). This packet is larger, though, as it contains the TLS session resumption information. Thus, it is a trade-off of a larger message versus a (possible) round trip savings. This message flow is shown below.

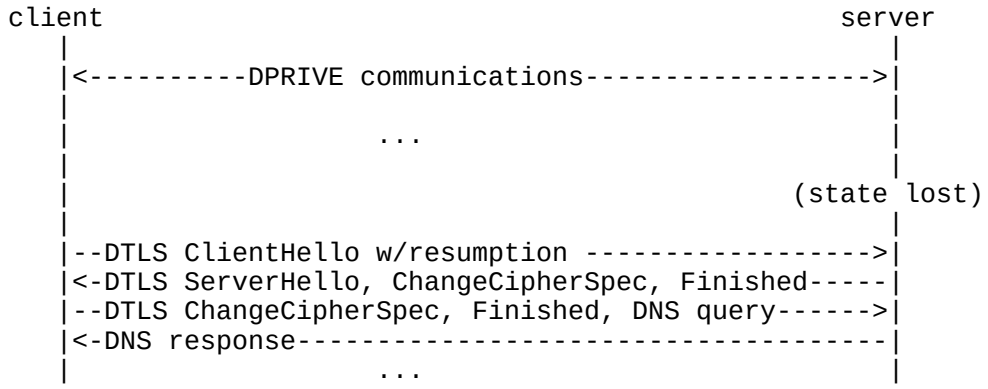

Figure 3: Server State Loss, DTLS False Start

2.3. TLS 1.3

Session resumption via identifiers and tickets is obsolete in TLS1.3 [I-D.ietf-tls-tls13]. Both methods are replaced by a pre-shared key (PSK) mode. A PSK is established on a previous connection after the handshake is completed, and can then be presented by the client on the next visit.

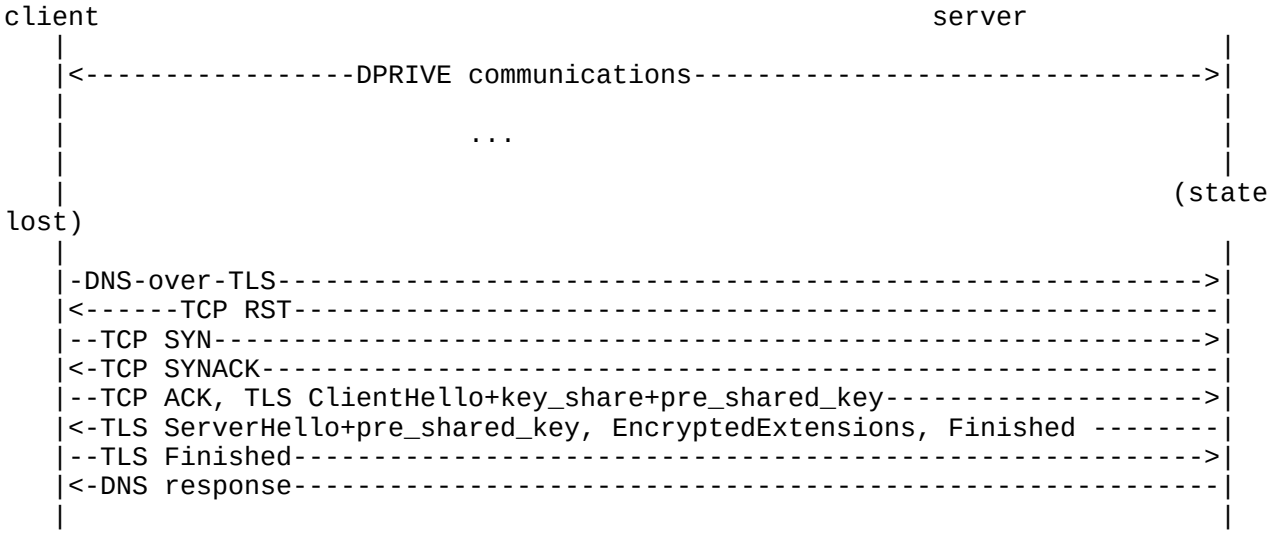

Figure 4: Session resumption

3. TCP Fast Open

If the client and server TCP stacks both support TCP Fast Open (TFO) [RFC7413], the TCP 3-way handshake can be done without a round trip, as shown below. Currently, TFO is supported in Linux 3.7 (TCP client and server), iOS 9, and OS X 10.11.

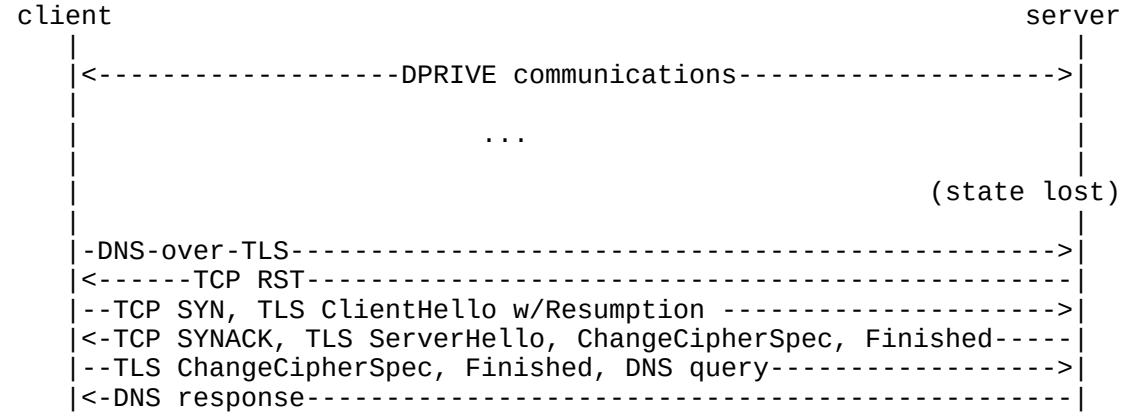

Figure 5: Server State Loss, TLS with TCP FastOpen

4. Probing for Server State Loss

In between normal DNS traffic while the communication to the DNS server is quiescent, the DNS client may want to probe the server to ensure it has maintained cryptographic state. Such probes can also keep alive firewall or NAT bindings. This probing reduces the frequency of needing a new handshake when a DNS query needs to be resolved, improving the user experience at the cost of bandwidth and processing time; cellular devices could even send the probes while in power-save state [I-D.isomaki-rtcweb-mobile].

If the server has lost state, a DTLS (or TLS) handshake needs to be initiated with the server.

4.1. DTLS

A DTLS heartbeat [RFC6520] verifies the server still has DTLS state by returning a DTLS message. If the server has lost state, it returns a DTLS Alert.

4.2. TLS

TLS runs over TCP, so a simple probe is a 0-length TCP packet (a "window probe"). This verifies the TCP connection is still working, which is also sufficient to prove the server has retained TLS state, because if the server loses TLS state it abandons the TCP connection. If the server has lost state, a TCP RST is returned immediately.

5. NAT or Firewall Pinhole Closed

A NAT or Firewall, on the path between the DPRIVE client and DPRIVE server, lose state -- either due to timing out the pinhole, exhausting its resources (and needing to prematurely close the pinhole), or crashing. This differs from the server losing state.

5.1. DTLS

With DTLS, the NAT or firewall will create a new mapping when it sees the new UDP packet. With a NAT, depending on its load (of other traffic) and its implmentation that mapping might be assigned to the same UDP port and IP address as the previous mapping, a different UDP port, and/or a different source IP address. The situation where the same mapping occurs is shown below.

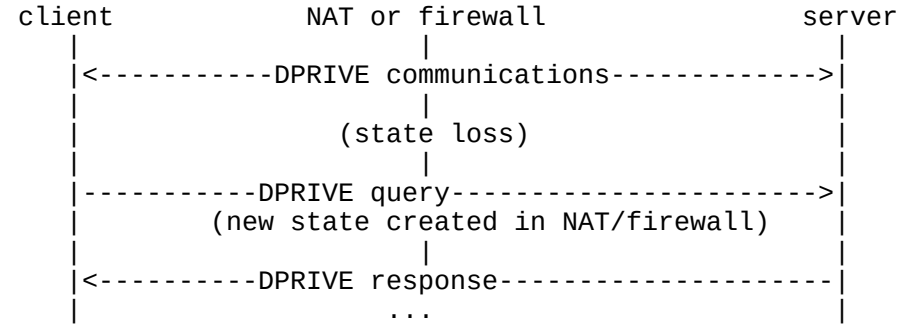

Figure 6: NAT/Firewall State Loss, DTLS

A different mapping can cause the server to reject the communication (DTLS Alert) cause the server to reject the communication (DTLS Alert) if the server was sensative to the client's source address or source port, consuming a round trip. This is shown below.

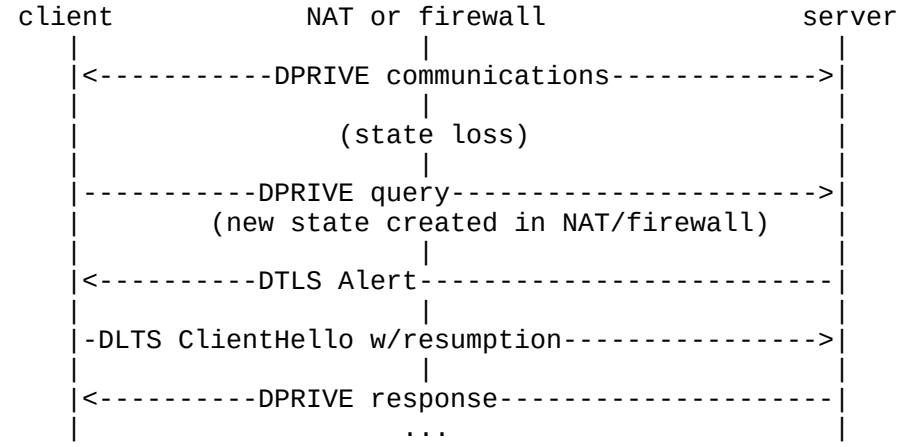

Figure 7: NAT/Firewall State Loss, DTLS, changed mapping

5.2. TLS

With a TCP connection when the NAT or firewall has lost state and sees a TCP packet without the SYN bit set, there are several possible reactions by the NAT or firewall:

- o send TCP RST, spoofing the source IP address of the original packet's destination address. This is shown in the following figure.

o create state. A firewall is unlikely to create state when it sees an in-progress TCP packet, but some NATs may create state. However, if the NAT creates state for a different source TCP port than the previous connection, the server will reject the TCP packet as shown in Figure 5.

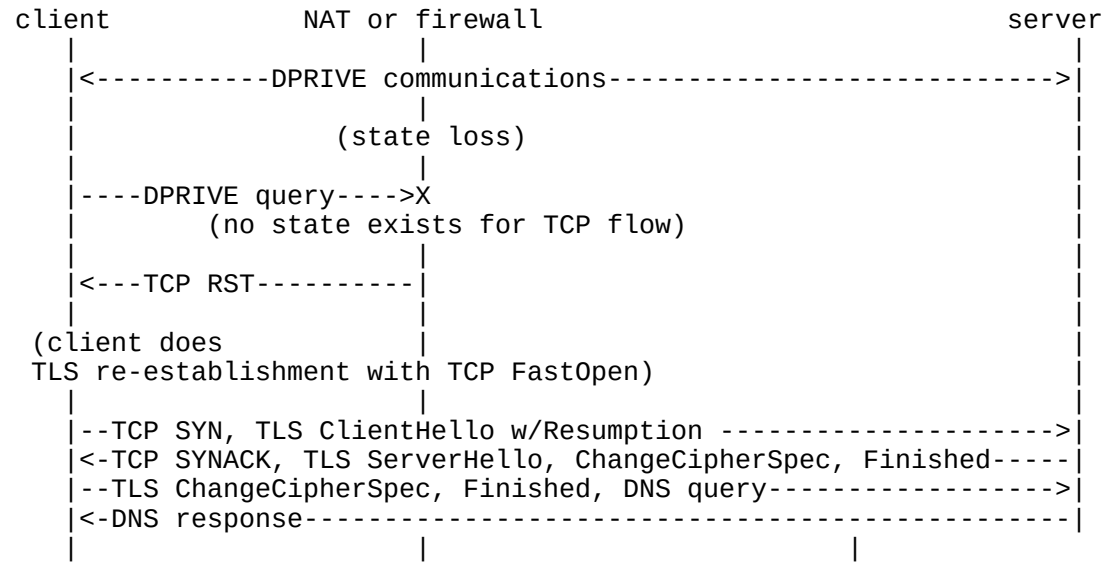

Figure 8: NAT/Firewall State Loss, TLS with TCP FastOpen

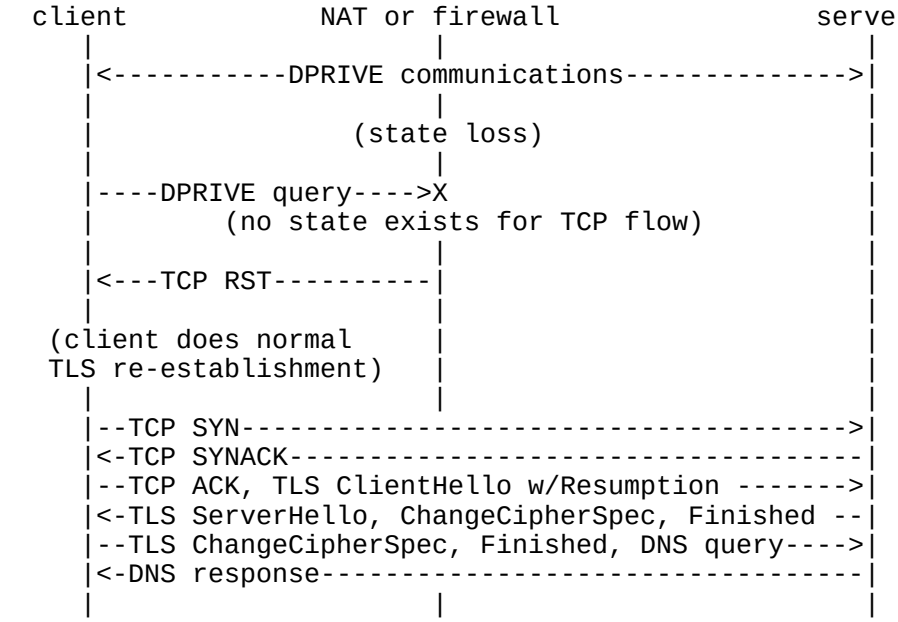

Figure 9: NAT/Firewall State Loss, TLS

6. IANA Considerations

None.

7. Acknowledgements

Authors would like to thank Allison Mankin for comments and review.

8. References

8.1. Normative References

[I-D.ietf-dprive-dns-over-tls]
Zi, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over TLS", draft-ietf-dprive-dns-over-tls-07 (work in progress), March 2016.

[I-D.ietf-dprive-dnsodtls]
Reddy, T., Wing, D., and P. Patil, "DNS over DTLS (DNSoD)", draft-ietf-dprive-dnsodtls-04 (work in progress), January 2016.

8.2. Informative References

[I-D.ietf-tls-tls13]
Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", draft-ietf-tls-tls13-11 (work in progress), December 2015.

[I-D.isomaki-rtcweb-mobile]
Isomaki, M., "RTCweb Considerations for Mobile Devices", draft-isomaki-rtcweb-mobile-00 (work in progress), July 2012.

[RFC6520] Seggelmann, R., Tuexen, M., and M. Williams, "Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) Heartbeat Extension", RFC 6520, DOI 10.17487/RFC6520, February 2012, <http://www.rfc-editor.org/info/rfc6520>.

[RFC7413] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, DOI 10.17487/RFC7413, December 2014, <http://www.rfc-editor.org/info/rfc7413>.

Authors' Addresses

Dan Wing
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134
USA

Email: dwing@cisco.com

Tirumaleswar Reddy
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tireddy@cisco.com