

Network  
Internet-Draft  
Intended status: Standards Track  
Expires: August 18, 2016

T. Pauly  
Apple Inc.  
S. Touati  
Ericsson  
R. Mantha  
Cisco Systems  
February 15, 2016

TCP Encapsulation of IKEv2 and IPSec Packets  
draft-pauly-ipsecme-tcp-encaps-03

## Abstract

This document describes a method to transport IKEv2 and IPSec packets over a TCP connection for traversing network middleboxes that may block IKEv2 negotiation over UDP. This method, referred to as TCP encapsulation, involves sending all packets for tunnel establishment as well as tunneled packets over a TCP connection. This method is intended to be used as a fallback option when IKE cannot be negotiated over UDP.

## 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 August 18, 2016.

## 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

## Internet-Draft TCP Encapsulation of IKEv2 and IPSec Packets February 2016

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">2</a>
<a href="#">1.1.</a>	<a href="#">Prior Work and Motivation . . . . .</a>	<a href="#">3</a>
<a href="#">1.2.</a>	<a href="#">Requirements Language . . . . .</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Configuration . . . . .</a>	<a href="#">4</a>
<a href="#">3.</a>	<a href="#">TCP-Encapsulated Header Formats . . . . .</a>	<a href="#">5</a>
<a href="#">3.1.</a>	<a href="#">TCP-Encapsulated IKEv2 Header Format . . . . .</a>	<a href="#">5</a>
<a href="#">3.2.</a>	<a href="#">TCP-Encapsulated ESP Header Format . . . . .</a>	<a href="#">6</a>
<a href="#">4.</a>	<a href="#">Applicability . . . . .</a>	<a href="#">6</a>
<a href="#">5.</a>	<a href="#">Connection Establishment and Teardown . . . . .</a>	<a href="#">6</a>
<a href="#">6.</a>	<a href="#">Interaction with NAT Detection Payloads . . . . .</a>	<a href="#">8</a>
<a href="#">7.</a>	<a href="#">Using MOBIKE with TCP encapsulation . . . . .</a>	<a href="#">8</a>
<a href="#">8.</a>	<a href="#">Using IKEv2 Message Fragmentation with TCP encapsulation . .</a>	<a href="#">8</a>
<a href="#">9.</a>	<a href="#">Considerations for Keep-alives and DPD . . . . .</a>	<a href="#">9</a>
<a href="#">10.</a>	<a href="#">Middlebox Considerations . . . . .</a>	<a href="#">9</a>
<a href="#">11.</a>	<a href="#">Performance Considerations . . . . .</a>	<a href="#">10</a>
<a href="#">11.1.</a>	<a href="#">TCP-in-TCP . . . . .</a>	<a href="#">10</a>
<a href="#">11.2.</a>	<a href="#">Added Reliability for Unreliable Protocols . . . . .</a>	<a href="#">10</a>
<a href="#">11.3.</a>	<a href="#">Encryption Overhead . . . . .</a>	<a href="#">10</a>
<a href="#">11.4.</a>	<a href="#">Quality of Service Markings . . . . .</a>	<a href="#">11</a>
<a href="#">11.5.</a>	<a href="#">Maximum Segment Size . . . . .</a>	<a href="#">11</a>
<a href="#">12.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">11</a>
<a href="#">13.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">11</a>
<a href="#">14.</a>	<a href="#">Acknowledgments . . . . .</a>	<a href="#">11</a>
<a href="#">15.</a>	<a href="#">References . . . . .</a>	<a href="#">12</a>
<a href="#">15.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">12</a>
<a href="#">15.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">12</a>
<a href="#">Appendix A.</a>	<a href="#">Example exchanges of TCP Encapsulation with TLS . .</a>	<a href="#">13</a>
<a href="#">A.1.</a>	<a href="#">Establishing an IKEv2 session . . . . .</a>	<a href="#">13</a>
<a href="#">A.2.</a>	<a href="#">Deleting an IKEv2 session . . . . .</a>	<a href="#">15</a>
<a href="#">A.3.</a>	<a href="#">Re-establishing an IKEv2 session . . . . .</a>	<a href="#">16</a>
<a href="#">A.4.</a>	<a href="#">Using MOBIKE between UDP and TCP Encapsulation . . . . .</a>	<a href="#">16</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">18</a>

[1.](#) Introduction

IKEv2 [[RFC7296](#)] is a protocol for establishing IPsec tunnels, using IKE messages over UDP for control traffic, and using Encapsulating Security Payload (ESP) messages for its data traffic. Many network middleboxes that filter traffic on public hotspots block all UDP traffic, including IKEv2 and IPsec, but allow TCP connections through

## Internet-Draft TCP Encapsulation of IKEv2 and IPsec Packets February 2016

since they appear to be web traffic. Devices on these networks that need to use IPsec (to access private enterprise networks, to route voice-over-IP calls to carrier networks, or because of security policies) are unable to establish IPsec tunnels. This document defines a method for encapsulating both the IKEv2 control messages as well as the IPsec data messages within a TCP connection.

Using TCP as a transport for IPsec packets adds a third option to the list of traditional IPsec transports:

1. Direct. Currently, IKEv2 negotiations begin over UDP port 500. If no NAT is detected between the initiator and the receiver, then subsequent IKEv2 packets are sent over UDP port 500 and IPsec data packets are sent using ESP [[RFC4303](#)].
2. UDP Encapsulation [[RFC3948](#)]. If a NAT is detected between the initiator and the receiver, then subsequent IKEv2 packets are sent over UDP port 4500 with four bytes of zero at the start of the UDP payload and ESP packets are sent out over UDP port 4500. Some peers default to using UDP encapsulation even when no NAT are detected on the path as some middleboxes do not support IP protocols other than TCP and UDP.
3. TCP Encapsulation. If both of the other two methods are not available or appropriate, both IKEv2 negotiation packets as well as ESP packets can be sent over a single TCP connection to the peer. This connection can itself use TLS [[RFC5246](#)] or other methods if needed. If the connection uses TLS, it will also be capable of traversing a web proxy [[RFC2817](#)].

Direct use of ESP or UDP Encapsulation SHOULD be preferred by IKEv2 implementations due to performance concerns when using TCP Encapsulation [Section 11](#). Most implementations should use TCP Encapsulation only on networks where negotiation over UDP has been attempted without receiving responses from the peer, or if a network is known to not support UDP.

## [1.1.](#) Prior Work and Motivation

Encapsulating IKEv2 connections within TCP or TLS streams is a common approach to solve the problem of UDP packets being blocked by network middleboxes. The goal of this document is to promote interoperability by providing a standard method of framing IKEv2 and ESP message within streams, and to provide guidelines for how to configure and use TCP encapsulation.

Some previous solutions include:

Pauly, et al.

Expires August 18, 2016

[Page 3]

---

Internet-Draft TCP Encapsulation of IKEv2 and IPSec Packets February 2016

Cellular Network Access Interworking Wireless LAN (IWLAN) uses IKEv2 to create secure connections to cellular carrier networks for making voice calls and accessing other network services over Wi-Fi networks. 3GPP has recommended that IKEv2 and ESP packets be sent within a TLS connection to be able to establish connections on restrictive networks.

ISAKMP over TCP Various non-standard extensions to ISAKMP have been deployed that send IPSec traffic over TCP or TCP-like packets.

SSL VPNs Many proprietary VPN solutions use a combination of TLS and IPSec in order to provide reliability.

IKEv2 over TCP IKEv2 over TCP as described in [\[I-D.nir-ipsecme-ike-tcp\]](#) is used to avoid UDP fragmentation.

## [1.2.](#) Requirements Language

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.](#) Configuration

One of the main reasons to use TCP encapsulation is that UDP traffic may be entirely blocked on a network. Because of this, support for TCP encapsulation is not specifically negotiated in the IKEv2 exchange. Instead, support for TCP encapsulation must be pre-configured on both the initiator and the responder.

The configuration defined on each peer should include the following parameters:

- o One or more TCP ports on which the responder will listen for incoming connections. Note that the initiator may initiate TCP connections to the responder from any local port.
- o Whether or not to use TLS for connections to a given TCP port. The responder may expect to read encapsulated IKEv2 and ESP packets directly from the TCP connection, or it may expect to read them from a stream of TLS data packets. The initiator should be pre-configured to use TLS or not when communicating with a given port on the responder.

This document leaves the selection of TCP ports up to implementations. If TLS is configured, many implementations will select to listen on TCP port 443 in order to traverse firewalls. If

TLS is not used, it is suggested to use TCP port 4500, which is allocated for IPSec NAT Traversal.

Since TCP encapsulation of IKEv2 and IPSec packets adds overhead and has potential performance trade-offs compared to direct or UDP-encapsulated tunnels (as described in Performance Considerations, [Section 11](#)), when possible, implementations SHOULD prefer IKEv2 direct or UDP encapsulated tunnels over TCP encapsulated tunnels.

### [3.](#) TCP-Encapsulated Header Formats

In order to encapsulate IKEv2 and ESP messages within a stream (which may be raw TCP, or TLS over TCP), a 32-bit length field precedes every message. If the first 32-bits of the message are zeros (a Non-ESP Marker), then the contents comprise an IKEv2 message. Otherwise, the contents comprise an ESP message. Authentication Header (AH) messages are not supported for TCP encapsulation.

Although a TCP stream may be able to send very long messages, implementations SHOULD limit message lengths to typical UDP datagram ESP payload lengths. The maximum message length is used as the effective MTU for connections that are being encrypted using ESP, so

the maximum message length will influence characteristics of inner connections, such as the TCP Maximum Segment Size (MSS).

### 3.1. TCP-Encapsulated IKEv2 Header Format


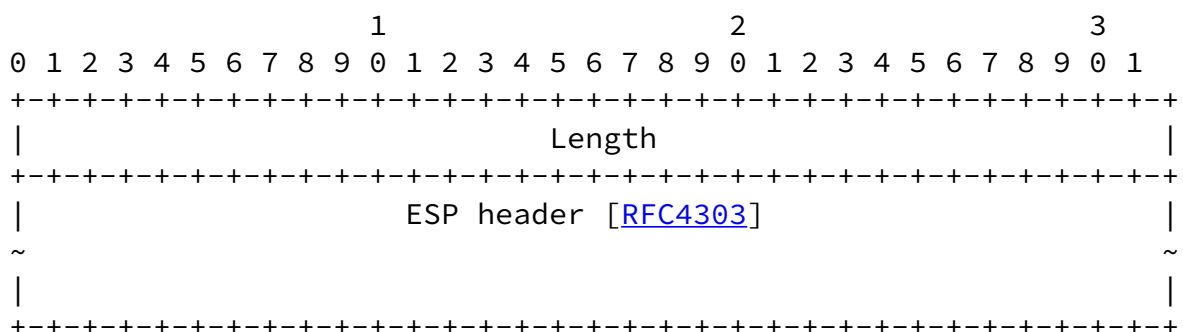


Figure 1

The IKE header is preceded by a 32-bit length field in network byte order that specifies the length of the IKE packet within the TCP stream. As with IKEv2 over UDP port 4500, a zeroed 32-bit Non-ESP Marker is inserted before the start of the IKEv2 header in order to differentiate the traffic from ESP traffic between the same addresses and ports.

- o Length (4 octets, unsigned integer) - Length of the IKE packet including the Length Field and Non-ESP Marker.

### 3.2. TCP-Encapsulated ESP Header Format



## Figure 2

The ESP header is preceded by a 32-bit length field in network byte order that specifies the length of the ESP packet within the TCP stream.

- o Length (4 octets, unsigned integer) - Length of the ESP packet including the Length Field.

### 4. Applicability

TCP encapsulation is applicable only when it has been configured to be used with specific IKEv2 peers. If a responder is configured to use TCP encapsulation, it **MUST** listen on the configured port(s) in case any peers will initiate new IKEv2 sessions. Initiators **MAY** use TCP encapsulation for any IKEv2 session to a peer that is configured to support TCP encapsulation, although it is recommended that initiators should only use TCP encapsulation when traffic over UDP is blocked.

Any specific IKE SA, along with its Child SAs, is either TCP encapsulated or not. A mix of TCP and UDP encapsulation for a single SA is not allowed. The exception to this rule is SAs that are migrated between addresses using MOBIKE [Section 7](#).

### 5. Connection Establishment and Teardown

When the IKEv2 initiator uses TCP encapsulation for its negotiation, it will initiate a TCP connection to the responder using the configured TCP port. If TLS is being used, it will be negotiated at this point. If a web proxy is applied to the ports for the TCP connection, and TLS is being used, the initiator can send an HTTP CONNECT message to establish a tunnel through the proxy [[RFC2817](#)].

Once this connection is established, encapsulated IKEv2 messages will negotiate the IKE SA and initial Child SA [[RFC7296](#)]. After this point, both encapsulated IKE Figure 1 and ESP Figure 2 messages will be sent over the TCP connection.

In order to close an IKE session, either the initiator or responder **SHOULD** gracefully tear down IKE SAs with DELETE payloads. Once all SAs have been deleted, the initiator of the original connection **MUST**

close the TCP connection.

An unexpected FIN or a RST on the TCP connection may indicate either a loss of connectivity, an attack, or some other error. If a DELETE payload has not been sent, both sides SHOULD maintain the state for their SAs for the standard lifetime or time-out period. The original initiator (that is, the endpoint that initiated the TCP connection and sent the first IKE\_SA\_INIT message) is responsible for re-establishing the TCP connection if it is torn down for any unexpected reason. Since new TCP connections may use different ports due to NAT mappings or local port allocations changing, the responder MUST allow packets for existing SAs to be received from new source ports.

If TLS is being used, TLS must be re-negotiated on any new TCP connections created due to a broken connection. TLS Session Resumption is recommended to improve efficiency.

A peer must discard a partially received message due to a broken connection.

The streams of data sent over any TCP connection used for this protocol MUST begin with a complete message, which is either an encapsulated IKE or ESP message. If the connection is being used to resume a previous IKE session, the responder can recognize the session using either the IKE SPI from an encapsulated IKE message or the ESP SPI from an encapsulated ESP message. If the session had been fully established previously, it is suggested that the initiator send an UPDATE\_SA\_ADDRESSES message if MOBIKE is supported, or an INFORMATIONAL message (a keepalive) otherwise. If either initiator or responder receives a stream that cannot be parsed correctly, it MUST close the TCP connection.

Multiple TCP connections between the initiator and the responder are allowed, but their use must take into account the initiator capabilities and the deployment model such as to connect to multiple gateways handling different ESP SAs when deployed in a high availability model. It is also possible to negotiate multiple IKE SAs over the same TCP connection.



single or multiple TCP connections.

## 6. Interaction with NAT Detection Payloads

When negotiating over UDP port 500, IKE\_SA\_INIT packets include NAT\_DETECTION\_SOURCE\_IP and NAT\_DETECTION\_DESTINATION\_IP payloads to determine if UDP encapsulation of IPSec packets should be used. These payloads contain SHA-1 digests of the SPIs, IP addresses, and ports. IKE\_SA\_INIT packets sent on a TCP connection SHOULD include these payloads, and SHOULD use the applicable TCP ports when creating and checking the SHA-1 digests.

If a NAT is detected due to the SHA-1 digests not matching the expected values, no change should be made for encapsulation of subsequent IKEv2 or ESP packets, since TCP encapsulation inherently supports NAT traversal. Implementations MAY use the information that a NAT is present to influence keep-alive timer values.

## 7. Using MOBIKE with TCP encapsulation

When an IKEv2 session is transitioned between networks using MOBIKE [[RFC4555](#)], the initiator of the transition may switch between using TCP encapsulation, UDP encapsulation, or no encapsulation. Implementations that implement both MOBIKE and TCP encapsulation MUST support dynamically enabling and disabling TCP encapsulation as interfaces change.

The encapsulation method of ESP packets MUST always match the encapsulation method of the IKEv2 negotiation, which may be different when an IKEv2 endpoint changes networks. When a MOBIKE-enabled initiator changes networks, the UPDATE\_SA\_ADDRESSES notification SHOULD be sent out first over UDP before attempting over TCP. If there is a response to the UPDATE\_SA\_ADDRESSES notification sent over UDP, then the ESP packets should be sent directly over IP or over UDP port 4500 (depending on if a NAT was detected), regardless of if a connection on a previous network was using TCP encapsulation. Similarly, if the responder only responds to the UPDATE\_SA\_ADDRESSES notification over TCP, then the ESP packets should be sent over the TCP connection, regardless of if a connection on a previous network did not use TCP encapsulation.

## 8. Using IKEv2 Message Fragmentation with TCP encapsulation

IKEv2 Message Fragmentation [[RFC7383](#)] is not required when using TCP encapsulation, since a TCP stream already handles the fragmentation of its contents across packets. Since fragmentation is redundant in this case, implementations might choose to not negotiate IKEv2

fragmentation. Even if fragmentation is negotiated, an implementation MAY choose to not fragment when going over a TCP connection.

If an implementation supports both MOBIKE and IKEv2 fragmentation, it SHOULD negotiate IKEv2 fragmentation over a TCP encapsulated session in case the session switches to UDP encapsulation on another network.

## 9. Considerations for Keep-alives and DPD

Encapsulating IKE and IPSec inside of a TCP connection can impact the strategy that implementations use to detect peer liveness and to maintain middlebox port mappings. Peer liveness should be checked using IKEv2 Informational packets [[RFC7296](#)].

In general, TCP port mappings are maintained by NATs longer than UDP port mappings, so IPSec ESP NAT keep-alives [[RFC3948](#)] SHOULD NOT be sent when using TCP encapsulation. Any implementation using TCP encapsulation MUST silently drop incoming NAT keep-alive packets, and not treat them as errors. NAT keep-alive packets over a TCP encapsulated IPSec connection will be sent with a length value of 1 byte, whose value is 0xFF Figure 2.

Note that depending on the configuration of TCP and TLS on the connection, TCP keep-alives [[RFC1122](#)] and TLS keep-alives [[RFC6520](#)] may be used. These MUST NOT be used as indications of IKEv2 peer liveness.

## 10. Middlebox Considerations

Many security networking devices such as Firewalls or Intrusion Prevention Systems, network optimization/acceleration devices and Network Address Translation (NAT) devices keep the state of sessions that traverse through them.

These devices commonly track the transport layer and/or the application layer data to drop traffic that is anomalous or malicious in nature.

A network device that monitors up to the application layer will commonly expect to see HTTP traffic within a TCP socket running over port 80, if non-HTTP traffic is seen (such as TCP encapsulated IKEv2), this could be dropped by the security device. In this scenario the IKEv2 and ESP packets are to be TLS encapsulated (using port 443) to ensure that the security device permits the traffic.

In the case of a TLS proxy, where the network security device decrypts, inspects and re-encrypts the session the same

## Internet-Draft TCP Encapsulation of IKEv2 and IPSec Packets February 2016

considerations must be taken into account as for HTTP, where TLS can be inspected. In this case the TLS proxy must allow TCP encapsulation of IKEv2 and IPsec otherwise the session could be dropped.

A network device that monitors the transport layer will track the state of TCP sessions, such as TCP sequence numbers. TCP encapsulation of IKEv2 should therefore use standard TCP behaviors to avoid being dropped by middleboxes.

### [11.](#) Performance Considerations

Several aspects of TCP encapsulation for IKEv2 and IPSec packets may negatively impact the performance of connections within the tunnel. Implementations should be aware of these and take these into consideration when determining when to use TCP encapsulation.

#### [11.1.](#) TCP-in-TCP

If the outer connection between IKEv2 peers is over TCP, inner TCP connections may suffer effects from using TCP within TCP. In particular, the inner TCP's round-trip-time estimation will be affected by the burstiness of the outer TCP. This will make loss-recovery of the inner TCP traffic less reactive and more prone to spurious retransmission timeouts.

#### [11.2.](#) Added Reliability for Unreliable Protocols

Since ESP is an unreliable protocol, transmitting ESP packets over a TCP connection will change the fundamental behavior of the packets. Some application-level protocols that prefer packet loss to delay (such as Voice over IP or other real-time protocols) may be negatively impacted if their packets are retransmitted by the TCP connection due to packet loss.

#### [11.3.](#) Encryption Overhead

If TLS or another encryption method is used on the TCP connection, there may be increased processing overhead for encrypting and

decrypting. This overhead may be experienced as a decrease in throughput on CPU-limited devices, or an increase in CPU usage or battery consumption on other devices, therefore the initiator and responder MUST allow the selection of NULL cipher when using TLS. Additionally, the TLS record introduces another layer of overhead, requiring more bytes to transmit a given IKEv2 and IPSec packet.

Internet-Draft TCP Encapsulation of IKEv2 and IPSec Packets February 2016

#### [11.4.](#) Quality of Service Markings

Quality of Service (QoS) markings, such as DSCP and Traffic Class, should be used with care on TCP connections used for encapsulation. Individual packets SHOULD NOT use different markings than the rest of the connection, since packets with different priorities may be routed differently and cause unnecessary delays in the connection.

#### [11.5.](#) Maximum Segment Size

A TCP connection used for IKEv2 encapsulation SHOULD negotiate its maximum segment size (MSS) in order to avoid unnecessary fragmentation of packets.

### [12.](#) Security Considerations

IKEv2 responders that support TCP encapsulation may become vulnerable to new Denial-of-Service (DoS) attacks that are specific to TCP, such as SYN-flooding attacks. Responders should be aware of this additional attack-surface.

Attackers may be able to disrupt the TCP connection by sending spurious RST packets. Due to this, implementations SHOULD make sure that IKE session state persists even if the underlying TCP connection is torn down.

If TLS is used on the encapsulating TCP connection it should not be considered as a security measure, and no TLS profile is recommended by this specification. The security of the IKEv2 session is entirely derived from the IKEv2 negotiation and key establishment.

### [13.](#) IANA Considerations

This memo includes no request to IANA.

TCP port 4500 is already allocated to IPSec. This port MAY be used for the protocol described in this document, but implementations MAY prefer to use other ports based on local policy. We foresee some implementations using TCP port 443 to more easily pass through some middleboxes [[I-D.tschofenig-hourglass](#)].

## [14.](#) Acknowledgments

The authors would like to acknowledge the input and advice of Stuart Cheshire, Delziel Fernandes, Yoav Nir, Christoph Paasch, Yaron Sheffer, David Schinazi, Graham Bartlett, Byju Pularikkal, March Wu and Kingwel Xie. Special thanks to Eric Kinnear for his implementation work.

Pauly, et al.

Expires August 18, 2016

[Page 11]

---

Internet-Draft TCP Encapsulation of IKEv2 and IPSec Packets February 2016

## [15.](#) References

### [15.1.](#) Normative References

- [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>>.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, [RFC 7296](#), DOI 10.17487/RFC7296, October 2014, <<http://www.rfc-editor.org/info/rfc7296>>.

### [15.2.](#) Informative References

- [I-D.nir-ipsecme-ike-tcp] Nir, Y., "A TCP transport for the Internet Key Exchange", [draft-nir-ipsecme-ike-tcp-01](#) (work in progress), July 2012.
- [I-D.tschofenig-hourglass] Tschofenig, H., "The New Waist of the Hourglass", [draft-tschofenig-hourglass-00](#) (work in progress), July 2012.

- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, [RFC 1122](#), DOI 10.17487/RFC1122, October 1989, <<http://www.rfc-editor.org/info/rfc1122>>.
- [RFC2817] Khare, R. and S. Lawrence, "Upgrading to TLS Within HTTP/1.1", [RFC 2817](#), DOI 10.17487/RFC2817, May 2000, <<http://www.rfc-editor.org/info/rfc2817>>.
- [RFC3948] Huttunen, A., Swander, B., Volpe, V., DiBurro, L., and M. Stenberg, "UDP Encapsulation of IPsec ESP Packets", [RFC 3948](#), DOI 10.17487/RFC3948, January 2005, <<http://www.rfc-editor.org/info/rfc3948>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), DOI 10.17487/RFC4303, December 2005, <<http://www.rfc-editor.org/info/rfc4303>>.
- [RFC4555] Eronen, P., "IKEv2 Mobility and Multihoming Protocol (MOBIKE)", [RFC 4555](#), DOI 10.17487/RFC4555, June 2006, <<http://www.rfc-editor.org/info/rfc4555>>.

---

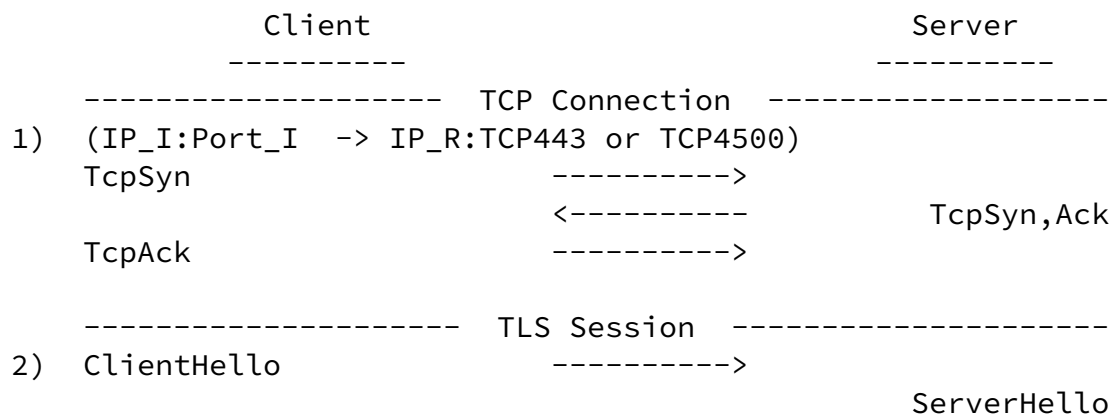
Internet-Draft TCP Encapsulation of IKEv2 and IPSec Packets February 2016

- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), DOI 10.17487/RFC5246, August 2008, <<http://www.rfc-editor.org/info/rfc5246>>.
- [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>>.
- [RFC7383] Smyslov, V., "Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation", [RFC 7383](#), DOI 10.17487/RFC7383, November 2014, <<http://www.rfc-editor.org/info/rfc7383>>.

[Appendix A](#). Example exchanges of TCP Encapsulation with TLS

## [A.1.](#) Establishing an IKEv2 session

Internet-Draft TCP Encapsulation of IKEv2 and IPSec Packets February 2016



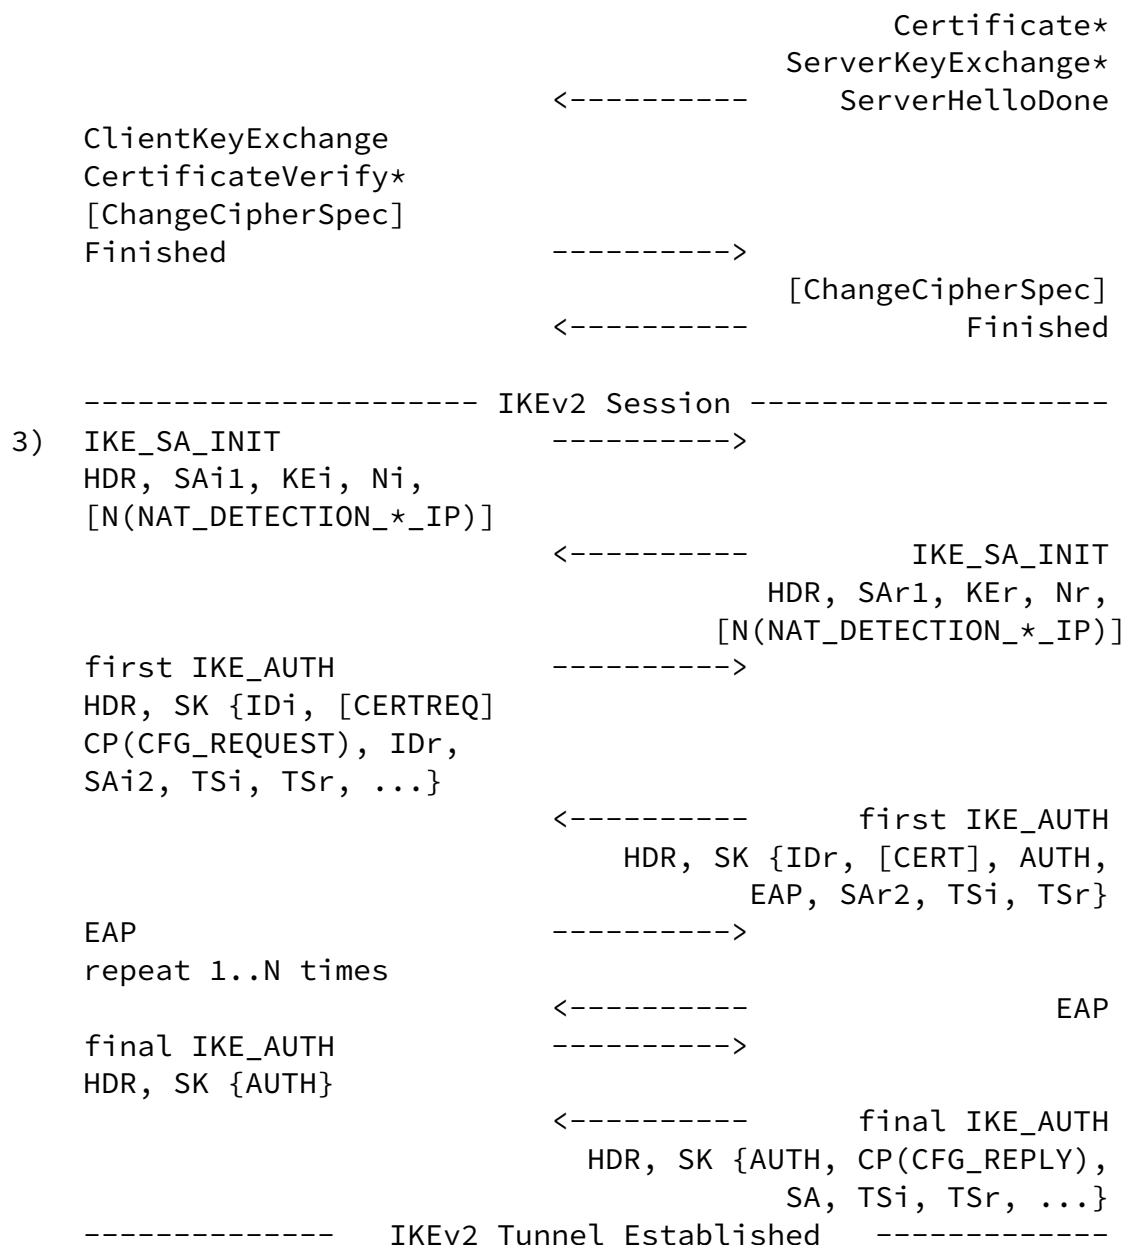


Figure 3

1. Client establishes a TCP connection with the server on port 443 or 4500.
2. Client initiates TLS handshake. During TLS handshake, the



server SHOULD NOT request the client's' certificate, since authentication is handled as part of IKEv2 negotiation.

3. Client and server establish an IKEv2 connection. This example shows EAP-based authentication, although any authentication type may be used.

#### [A.2.](#) Deleting an IKEv2 session

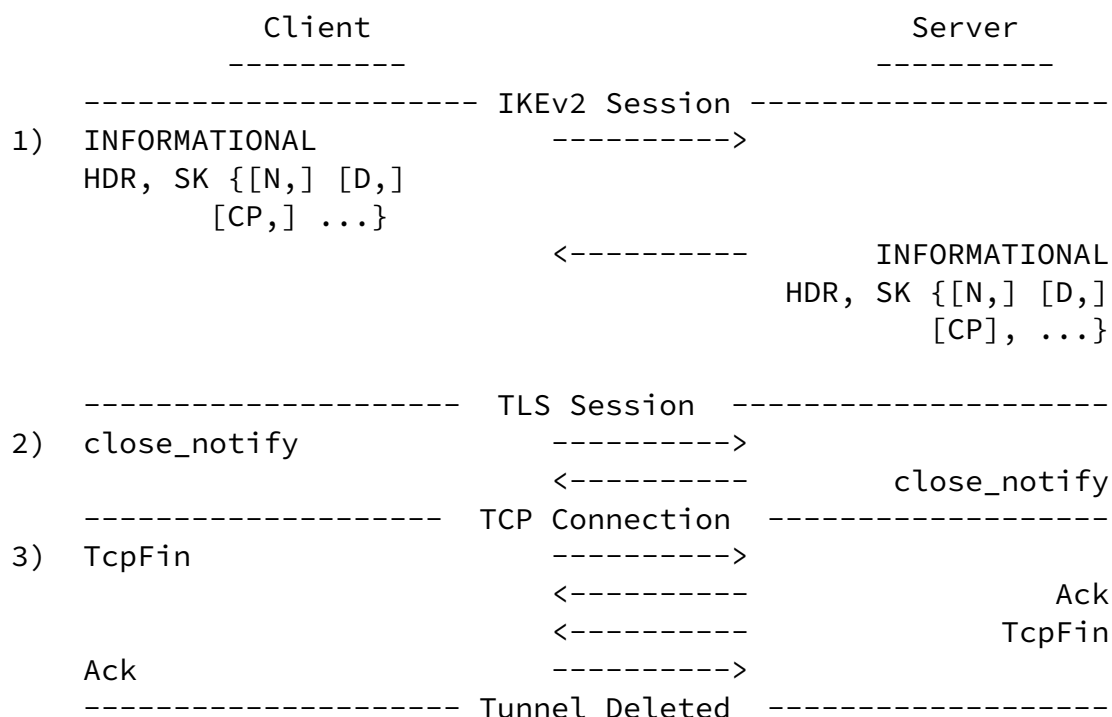


Figure 4

1. Client and server exchange INFORMATIONAL messages to notify IKE SA deletion.
2. Client and server negotiate TLS session deletion using TLS CLOSE\_NOTIFY.
3. The TCP connection is torn down.

Unless the TCP connection and/or TLS session are being used for multiple IKE SAs, the deletion of the IKE SA should lead to the disposal of the underlying TLS and TCP state.

### A.3. Re-establishing an IKEv2 session

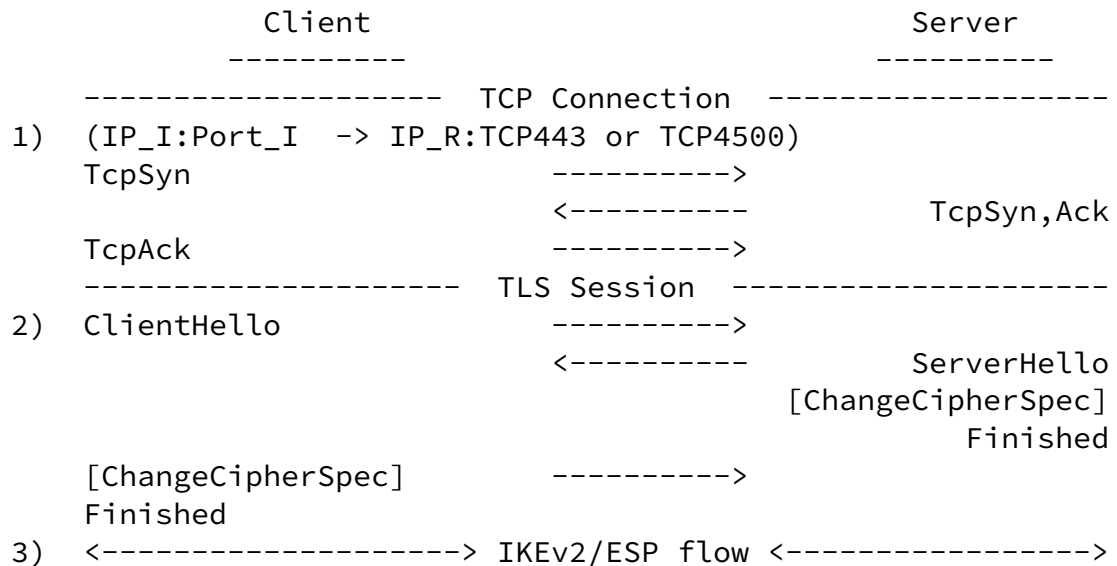
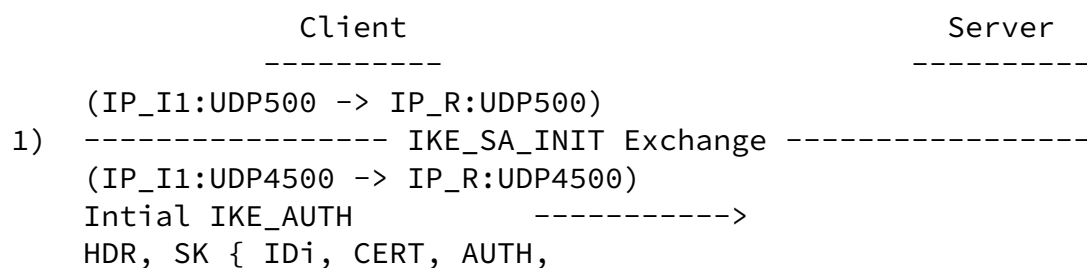


Figure 5

1. If a previous TCP connection was broken (for example, due to a RST), the client is responsible for re-initiating the TCP connection. The initiator's address and port (IP\_I and Port\_I) may be different from the previous connection's address and port.
2. In ClientHello TLS message, the client SHOULD send the Session ID it received in the previous TLS handshake if available. It is up to the server to perform either an abbreviated handshake or full handshake based on the session ID match.
3. After TCP and TLS are complete, the IKEv2 and ESP packet flow can resume. If MOBIKE is being used, the initiator SHOULD send UPDATE\_SA\_ADDRESSES.

### A.4. Using MOBIKE between UDP and TCP Encapsulation



CP(CFG\_REQUEST),

---

Internet-Draft TCP Encapsulation of IKEv2 and IPSec Packets February 2016

```
SAi2, TSi, TSr,
N(MOBIKE_SUPPORTED) }

<----- Initial IKE_AUTH
HDR, SK { IDr, CERT, AUTH,
          EAP, SAr2, TSi, TSr,
          N(MOBIKE_SUPPORTED) }
<----- IKEv2 tunnel establishment ----->

2) ----- MOBIKE Attempt on new network -----
(IP_I2:UDP4500 -> IP_R:UDP4500)
INFORMATIONAL ----->
HDR, SK { N(UPDATE_SA_ADDRESSES),
N(NAT_DETECTION_SOURCE_IP),
N(NAT_DETECTION_DESTINATION_IP) }

3) ----- TCP Connection -----
(IP_I2:PORT_I -> IP_R:TCP443 or TCP4500)
TcpSyn ----->
<----- TcpSyn,Ack
TcpAck ----->

4) ----- TLS Session -----
ClientHello ----->
ServerHello
Certificate*
ServerKeyExchange*
<----- ServerHelloDone
ClientKeyExchange
CertificateVerify*
[ChangeCipherSpec]
Finished ----->
[ChangeCipherSpec]
<----- Finished

5) ----- IKEv2 Session -----
INFORMATIONAL ----->
HDR, SK { N(UPDATE_SA_ADDRESSES),
N(NAT_DETECTION_SOURCE_IP),
N(NAT_DETECTION_DESTINATION_IP) }
```

```

                                <----- INFORMATIONAL
                                HDR, SK { N(NAT_DETECTION_SOURCE_IP),
                                N(NAT_DETECTION_DESTINATION_IP) }
6) <----- IKEv2/ESP data flow ----->

```

Figure 6

## Internet-Draft TCP Encapsulation of IKEv2 and IPSec Packets February 2016

1. During the IKE\_SA\_INIT exchange, the client and server exchange MOBIKE\_SUPPORTED notify payloads to indicate support for MOBIKE.
2. The client changes its point of attachment to the network, and receives a new IP address. The client attempts to re-establish the IKEv2 session using the UPDATE\_SA\_ADDRESSES notify payload, but the server does not respond because the network blocks UDP traffic.
3. The client beings up a TCP connection to the server in order to use TCP encapsulation.
4. The client initiates and TLS handshake with the server.
5. The client sends the UPDATE\_SA\_ADDRESSES notify payload on the TCP encapsulated connection.
6. The IKEv2 and ESP packet flow can resume.

### Authors' Addresses

Tommy Pauly  
 Apple Inc.  
 1 Infinite Loop  
 Cupertino, California 95014  
 US

Email: [tpauly@apple.com](mailto:tpauly@apple.com)

Samy Touati  
 Ericsson

300 Holger Way  
San Jose, California 95134  
US

Email: [samy.touati@ericsson.com](mailto:samy.touati@ericsson.com)

Ravi Mantha  
Cisco Systems  
SEZ, Embassy Tech Village  
Panathur, Bangalore 560 037  
India

Email: [ramantha@cisco.com](mailto:ramantha@cisco.com)