TCP Encapsulation of IKEv2 and IPSec Packets

Tommy Pauly (tpauly@apple.com)
Samy Touati (samy.touati@ericsson.com)
Ravi Mantha (ramantha@cisco.com)

IPSECME
IETF 95, April 2016, Buenos Aires

Status

- On third revision: https://tools.ietf.org/html/draft-pauly-ipsecme-tcp-encaps-03
  - Added clarifying details and an appendix of example packet exchanges
- IKEv2 over TCP has been proposed as an item the revised IPSecME charter
- Please review and give feedback!

Motivation

- TCP encapsulation of IKEv2 is necessary to traverse networks that block UDP
- There is no standard way to use IKEv2 over TCP
  - 3GPP spec[1] recommends using TCP/TLS, without protocol details
  - IKEv1 had proprietary implementations over TCP
  - "SSL" VPNs use TCP, but are not standardized

[1] 3GPP TS 33.402 version 12.5.0 Release 12

Goals

- Standardize framing headers for IKE and ESP messages within a stream
- Define minimum required configuration required by peers (TCP port, whether or not to use TLS)
- Provide guidelines for IKE session management when using TCP
  - How to handle TCP connection errors
  - MOBIKE

Non-Goals

- Avoid all firewalls that try to block tunnels
- Use a mix of TCP and UDP for the same IKE session simultaneously
- Use TCP to avoid fragmentation
- Define explicit algorithm for how implementations should choose UDP vs. TCP vs. TLS/TCP

Performance Considerations

- TCP retransmissions on the tunnel connection can cause performance problems in high-loss networks
  - Experiments showed good performance up to 3% loss
  - Excessive loss can cause packets to be received in bursts
    - If the tunnel's TCP connection has a small congestion window, this can be amplified for inner TCP streams
- Overhead of TCP and TLS headers increases amount of overall data
- We recommend TCP as only a fallback from UDP

Framing formats

3.1. TCP-Encapsulated IKEv2 Header Format

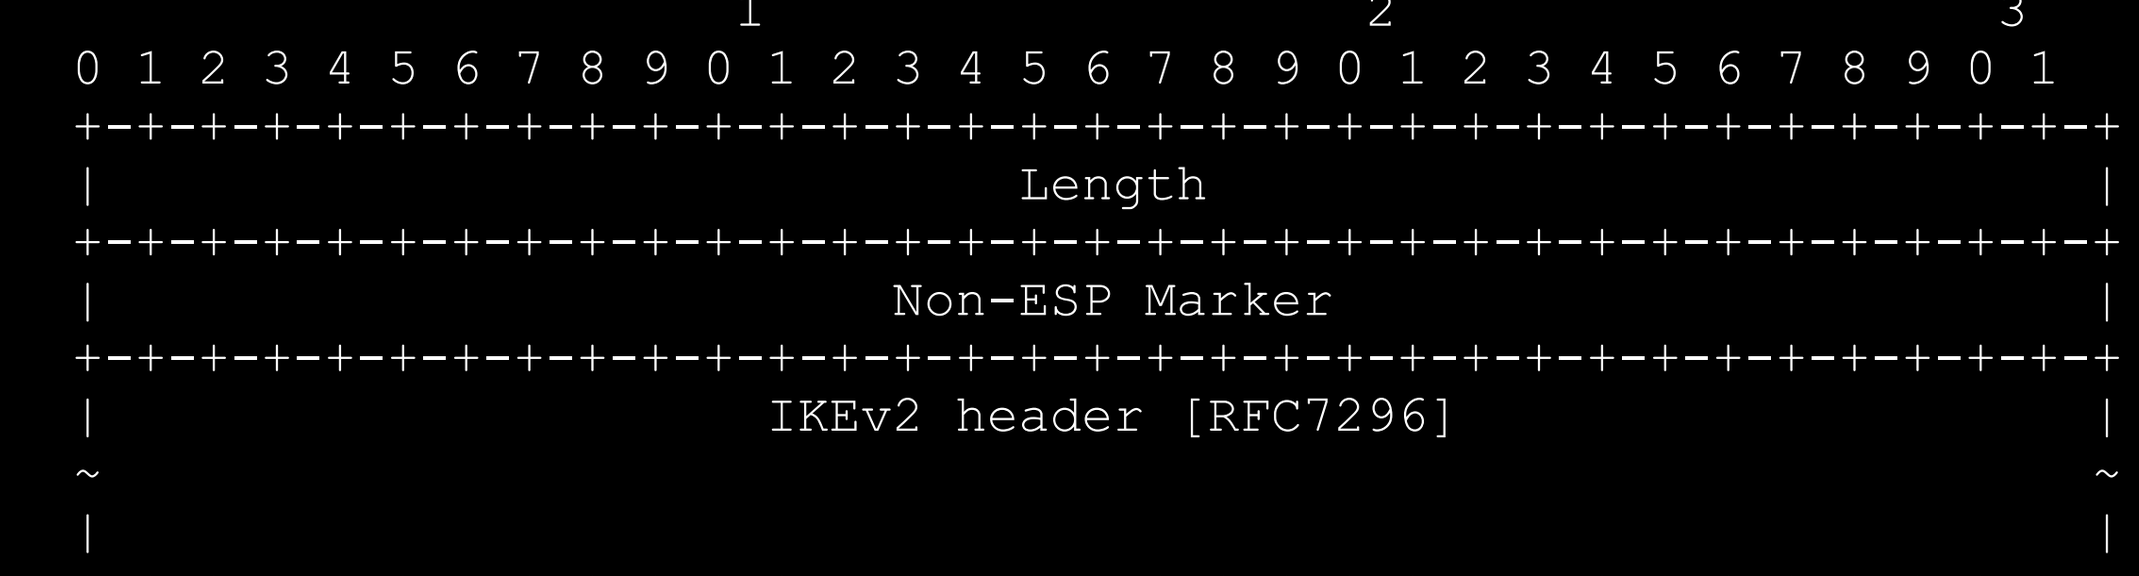

3.2. TCP-Encapsulated ESP Header Format

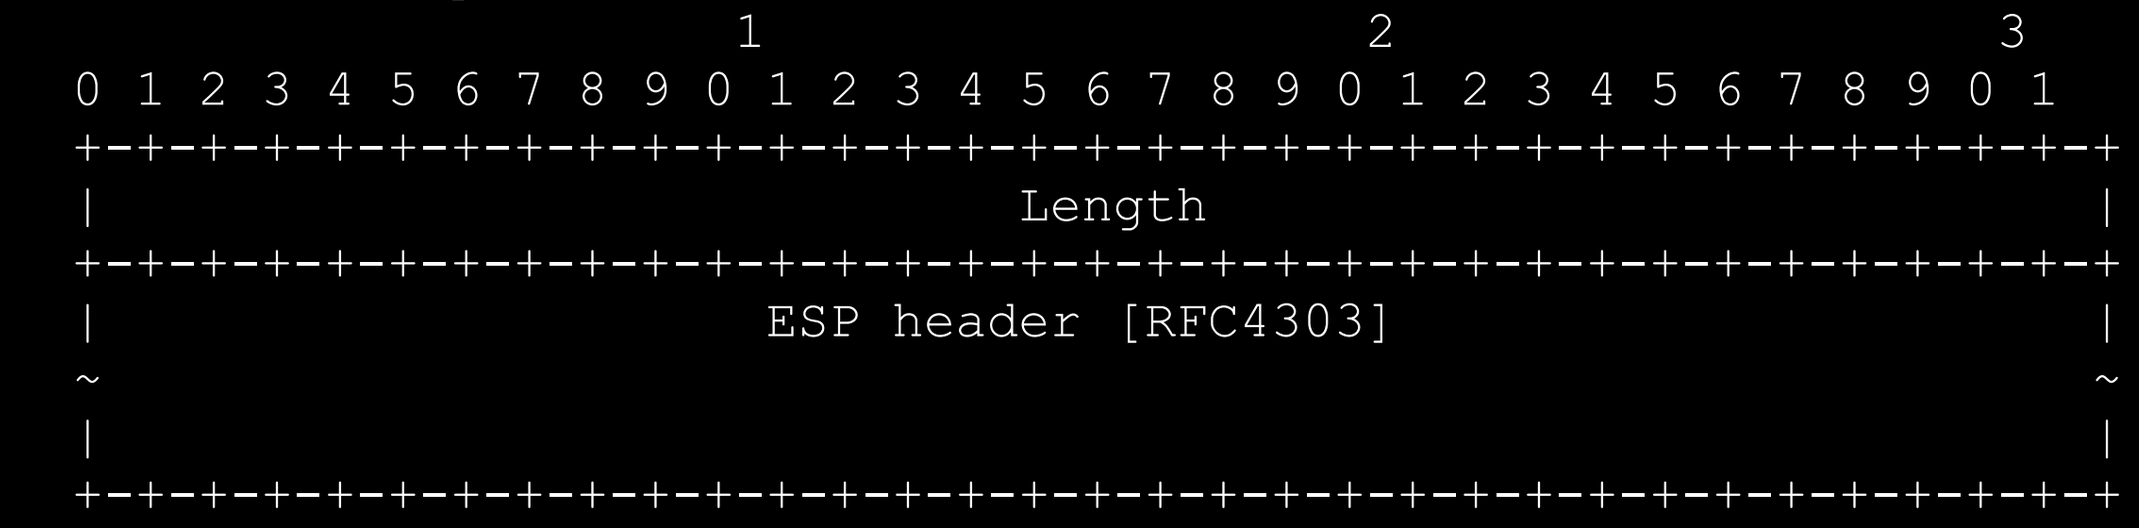

Session Lifetime

- IKEv2 session lifetime is completely independent of TCP connections or TLS sessions: connections being reset should not delete the IKE SA
- Using multiple TCP connections between peers is possible, as is multiple IKE SAs over a single TCP connection
- Allowed to switch to or from TCP encapsulation when using MOBIKE

Usage of TLS

- TLS (or use of TCP port 443 at least) is required to traverse certain networks, and will certainly be used by many implementations
- The framing protocol for TCP encapsulation does not need to change if the negotiation occurs over a TLS-encrypted stream
- Draft does not specify which port to use, but explains how the TLS session interacts with the IKE session, and how TLS authentication should not have any effect on IKE authentication
- Is this an appropriate level of detail regarding TLS?

Next steps

- Create an -04 update of the draft based on recent feedback
- Consider a secondary Informational Draft to explain deployment details: algorithm for choosing TCP after failing UDP; port selection; TLS details; and interactions with proxies
- Request for working group support and adoption