Workgroup: IPSECME Working Group
Internet-Draft:
draft-pan-ipsecme-esp-trailer-adjustment-00
Published: 23 October 2023
Intended Status: Informational
Expires: 25 April 2024
Authors: W. Pan C. Fang
Huawei Huawei

Considerations for Adjustments of Encapsulating Security Payload (ESP) Trailer

Abstract

Implementing IPsec in hardware is a way to improve the forwarding performance of IPsec. However, the current IPsec ESP packet design, i.e., the position of ESP trailer, imposes a new overhead challenge for implementing IPsec in hardware. This document explains how this overhead challenge occurs and proposes the possible solutions.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-pan-ipsecme-esp-trailer-adjustment/.

Discussion of this document takes place on the ipsec Working Group mailing list (mailto:ipsec@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/ipsec/. Subscribe at https://www.ietf.org/mailman/listinfo/ipsec/.

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 https://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 25 April 2024.

Table of Contents

1. Introduction

IPsec [RFC4301] is widely used to provide security for IP communications. It can adapt to various scenarios with the two protocols of AH [RFC4302] and ESP [RFC4303] and the two modes of tunnel and transport. For example, IPsec can create a secure tunnel between two gateways to protect traffic between sites, such as between an enterprise's branch and its headquarters and between data centers.

IPsec has an impact on forwarding performance due to the increased overhead of encryption and decryption processing of IP packets. Hence, the performance improvement of IPsec has been a significant concern in the industry, especially in the above scenarios where the original traffic rate is considerably high.

There have been some different approaches to address the performance issues of IPsec. One is to use more efficient cryptographic algorithms. For example, using the AES-GCM algorithm [RFC4106] provides the same level of security as the AES-CBC algorithm but is more efficient. Another is to use cryptographic hardware

acceleration. Hardware accelerators use specialized hardware to perform encryption and decryption operations (other operations of IPsec remain being implemented by software).

MACsec [MACsec] is a Layer 2 security protocol that provides security for networks like data centers and enterprise networks. Since the requirement of hop-by-hop deployment, MACsec is more complex to deploy when compared to IPsec, but it has the distinct advantage of high forwarding performance, which can even reach the line rate. One important reason is that the entire MACsec process (not just encryption and decryption processing) is implemented by hardware. Similarly, utilizing hardware to implement the whole IPsec process is also an effective means to improve IPsec performance.

However, the position of the ESP trailer in the IPsec ESP-encapsulated packets poses a new overhead challenge for implementing IPsec in hardware.

2. Problem Statement

Implementing IPsec in hardware is an ideal way to improve performance, but it is also expensive. Not to mention the cost of designing and manufacturing the hardware, IPsec ESP imposes additional costs on the hardware due to its protocol design.

The IPsec ESP protocol places the Next Header and Padding in the ESP trailer at the end of the packet. Based on the current design, after processing the plain text of the IPsec ESP packet (ESP Header and the content before), the (receiving) hardware can't use the manner of "decrypting and then transmitting" to deal with the Payload Data. It must cache the decrypted payload because it doesn't know how to reorganize the packet header before the payload, until it gets the Next Header and Padding information from the ESP trailer. For example, in the ESP tunnel mode, the EtherType field in the Ethernet frame should be set according to whether the inner packet is IPv4 or IPv6. In another example, in the ESP transport mode, the length field in the IP header should be adjusted according to the Padding length.

Such a requirement for caching payload is at a considerable cost to the hardware chip, including increased chip area and packet processing latency. A larger chip area means greater power consumption, which is not eco-friendly and not in line with the current trend towards green energy efficiency.

3. Possible Solution

Suppose the Next Header and Padding information can be obtained immediately after processing the ESP Header, even in the cipher text. In that case, the hardware can perform all operations

sequentially without additional caching. The possible solutions to accomplish this goal are listed below.

3.1. Simplified Next Header Judgment in ESP Tunnel Mode

In ESP tunnel mode, the inner payload is usually an IPv4 or IPv6 packet, and the first byte of both types of packets indicates the IP protocol version (4 for IPv4, 6 for IPv6). Therefore, instead of determining the protocol type of inner payload based on the Next Header field in the ESP trailer, one possible solution is to decrypt the first byte after the ESP Header and determine whether the inner payload is IPv4 or IPv6 based on the value of that byte.

However, the ESP protocol [RFC4303] describes the behavior of sending and discarding the dummy packet in support of traffic flow confidentiality. The dummy packet is generated by the sender by randomly generating a payload and discarded by the receiver by recognizing the value of 59 in the Next Header field. Since the payload is randomly generated, the value of its first byte maybe 4 or 6, which will cause the receiver to misjudge when this simplified Next Header judgment method is used.

In an enterprise or data center network scenario, the two gateways usually do not send the dummy packet after establishing an ESP tunnel. Therefore, this simplified judgment is a relatively practical solution in these scenarios. To be on the safe side, adding a new negotiation between the two gateways can be considered, i.e., to negotiate not sending dummy packets when creating a Child SA in IKEv2 [RFC7296].

3.2. Adjust ESP Packet Format

Another solution is to change the ESP packet format by moving the ESP trailer's position to right after the ESP Header, still retaining encryption on the ESP trailer.

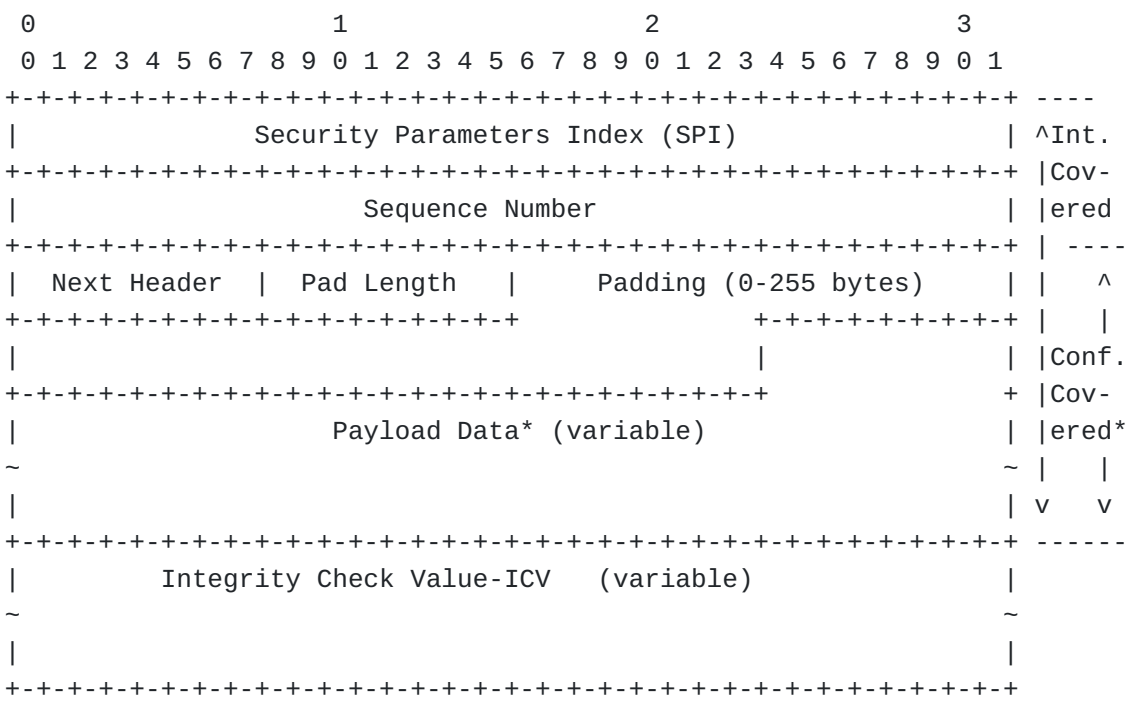

Figure 1: New ESP Packet Format

For the consideration of compatibility, the IPsec peers need to explicitly negotiate the use of this new ESP format when creating the Child SA.

This solution works not only for ESP tunnel mode but also for ESP transport mode. However, accordingly, it has more changes to the protocol.

4. Security Considerations

TBD

5. IANA Considerations

TBD

6. References

6.1. Normative References

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <https://www.rfc-editor.org/rfc/rfc4303>.

[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, <https://www.rfc-editor.org/rfc/rfc7296>.

6.2. Informative References

[MACsec] IEEE, "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Security", IEEE Std 802.1AE-2018 , 27 September 2018, <https://ieeexplore.ieee.org/document/8585421>.

[RFC4106] Viega, J. and D. McGrew, "The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)", RFC 4106, DOI 10.17487/RFC4106, June 2005, <https://www.rfc-editor.org/rfc/rfc4106>.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <https://www.rfc-editor.org/rfc/rfc4301>.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <https://www.rfc-editor.org/rfc/rfc4302>.

Authors' Addresses

Wei Pan
Huawei Technologies
China

Email: william.panwei@huawei.com

Chenyuan Fang
Huawei Technologies
China

Email: fangchenyuan@hisilicon.com