

IP Security Maintenance and Extensions
Internet-Draft
Intended status: Standards Track
Expires: 4 June 2026

C. Li
G. Zeng, Ed.
M. Zhang
X. Ding
Huawei Technologies
1 December 2025

IP Payload Compression excluding transport layer
draft-ls-ipsecme-ipcomp-exclude-transport-layer-02

Abstract

IP Payload Compression Protocol (IPComp) is used for compressing the IP payload in transmission to increase communication performance. The IPComp is applied to the payload of the IP datagram, starting with the first octet immediately after the IP header in IPv4, and the first octet after the excluded IPv6 Extension headers. However, transport layer information such as source port and destination port are useful in many network functions in transmission.

This document defines extensions of IP payload compression protocol (IPComp) to support compressing the payload excluding the transport layer information, to enable network functions using transport layer information (e.g., ECMP) working together with the payload compression. This document also defines an extension of IPComp to indicate the payload is not compressed to solve the out-of-order problems between the compressed and uncompressed packets.

About This Document

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

The latest revision of this draft can be found at <https://VMatrix1900.github.io/ipcomp-exclude-transport-layer/draft-ls-6man-ipcomp-exclude-transport-layer.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ls-ipsecme-ipcomp-exclude-transport-layer/>.

Discussion of this document takes place on the IP Security Maintenance and Extensions 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/>.

Source for this draft and an issue tracker can be found at <https://github.com/VMatrix1900/ipcomp-exclude-transport-layer>.

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 4 June 2026.

Copyright Notice

Copyright (c) 2025 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
2.1. Requirements Language	4
3. Problem Statement	4
4. Extensions to IPComp	4
4.1. Four-bytes Exclusion Extension	4
4.2. Uncompressed Payload Extension	6
5. IANA Considerations	7
6. Security Considerations	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Appendix A. Contributors	8
Authors' Addresses	8

1. Introduction

The IP Payload Compression Protocol (IPComp) [RFC3173] is defined to compress the IP payload in transmission in order to increase the communication performance between a pair of communicating nodes, provided the nodes have sufficient computation power and the communication is over slow or congested links.

In IP version 4, the compression is applied to the payload of the IP datagram, starting at the first octet following the IP header, and continuing through the last octet of the datagram. In the IPv6 context, IPComp is viewed as an end-to-end payload, and is not applied to IPv6 extension headers such as hop-by-hop, routing, and fragmentation extension headers[RFC8200]. The compression is applied starting at the first IP Header Option field that does not carry information that must be examined and processed by nodes along a packet's delivery path, if such an IP Header Option field exists, and continues to the ULP payload of the IP datagram. Therefore, the transport layer information such as source port and destination port is compressed. When IPComp is used, the Next Header field of IP header is set to 108, IPComp Datagram. The IPComp header contains the original Next Header and the Compress Parameter Index(CPI) is inserted between the IP header and the compressed payload.

There are many network functions which needs the transport layer information to work. For example, flow-based ECMP, Carrier Grade Network Translation (CGNAT), Access Control List (ACL) may require source and destination port to identify the transport layer flow. Some Firewall (FW), Deep Packet Inspection (DPI) also need to inspect the transport layer information. If IPComp compressed those transport layer information, the nodes along the packet's delivery path can not obtain the source port and destination port. Therefore the IPComp is not compatible with the network functions requiring the transport layer information which makes it harder to deploy.

This document defines an extension of IPComp to support compressing the payload excluding the first 4 bytes of transport layer header which contains source port and destination port. In this way, the IPComp can coexist with many network functions which requires these information. This document also defines an extension to explicitly indicate the payload is uncompressed to solve the out-of-order processing between the compressed and uncompressed packets.

2. Terminology

This document leverages the terms defined in [RFC3173]. The reader is assumed to be familiar with this terminology.

2.1. Requirements Language

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Problem Statement

Currently, the IPComp will compress all the IP payload which includes the transport layer information. If a layer 4 load balancer is deployed along the IPComp packet delivery path, then the load balancer can not obtain the source port and destination port to identify a flow without decompressing it first. In other words, the network functions which requires the transport layer information would also need to act as the decompression node of IPComp. This incompatibility makes the deployment of IPComp harder.

4. Extensions to IPComp

This section defines two extensions of IPComp. The first extension is used to indicate the first four bytes of transport layer header which contains the source port and destination is excluded from the compression. The second extension indicates that the payload is not compressed.

4.1. Four-bytes Exclusion Extension

This extension is used to indicate that the first four bytes of the transport layer header is excluded from the compression. The packet format using this extension is shown in Figure 1 (Demonstrated using IPv6 packet):

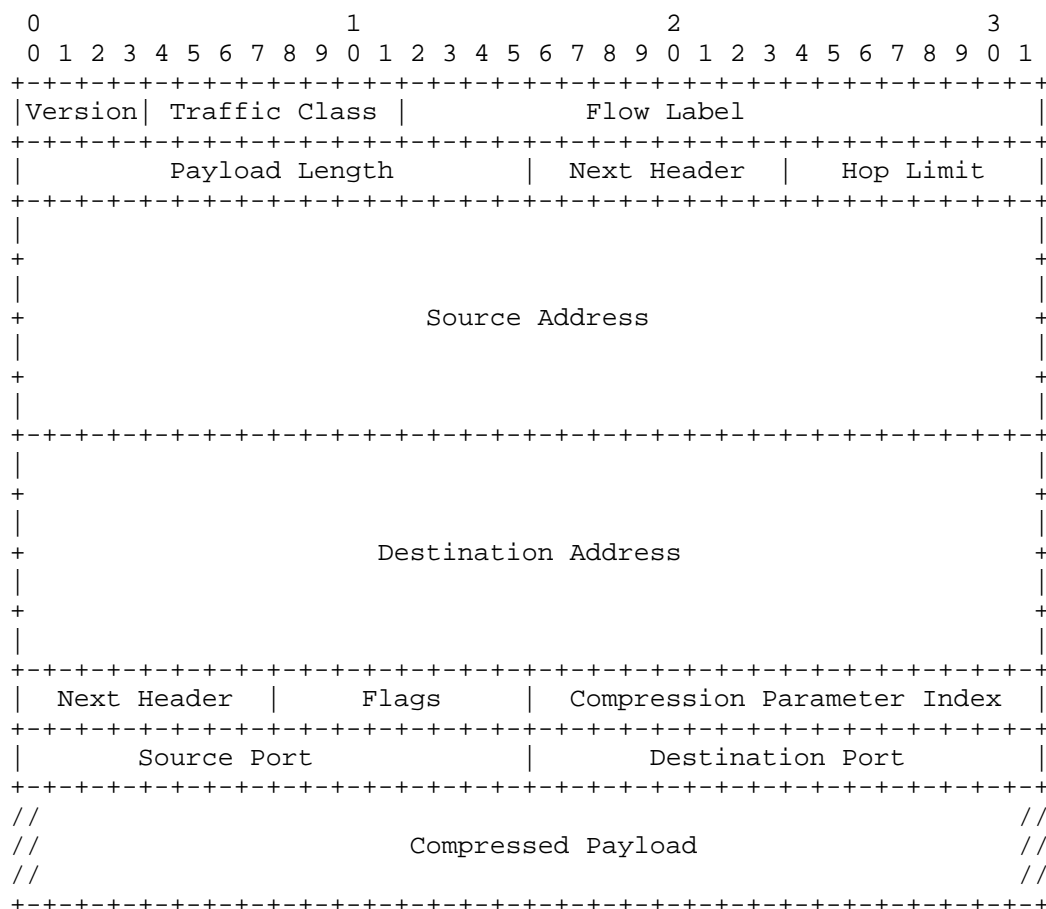


Figure 1: Packet format when using Four-bytes Exclusion Extension

To accomplish that there are two options to extend IPComp. The first option is to change the CPI field. Currently the CPI field identifies a particular compression algorithm. The defined CPI value can be found at [CPI-IANA]. We can define new CPI values to indicate the same compression algorithm with different compression range as shown in Table 1.

Value	Transform ID	References
0	RESERVED	[RFC2407]
1	IPCOMP_OUI	[RFC2407]
2	IPCOMP_DEFLATE	[RFC2407]
3	IPCOMP_LZS	[RFC2407]
4	IPCOMP_LZJH	[RFC3051]
TBD	IPCOMP_OUI with four bytes exclusion	This document
TBD	IPCOMP_DEFLATE with four bytes exclusion	This document
TBD	IPCOMP_LZS with four bytes exclusion	This document
TBD	IPCOMP_LZJH with four bytes exclusion	This document

Table 1: CPI with exclusion range Registry Entries

The second option is to change the Flags field. Currently, the Flags field is zero and ignored by the receiving node. We can introduce a bit to indicate whether the first four bytes is excluded from the compression range or not.

Which option is more suitable will be determined based on the discussion in the working group.

4.2. Uncompressed Payload Extension

Currently, if the total size of a compressed payload and the IPComp header is not smaller than the size of the original payload, the IP datagram will be sent in the original non-compressed form without the IPComp header. In the receiving node, the packet with the IPComp header will go through the decompression co-processor first while the packet without the IPComp header will be forwarded directly. Going through different packet process path will cause the out-of-order of packets within the same flow, reducing the transport performance.

To solve the out-of-order packets within the same IPComp-enabled flow, we propose to add IPComp header no matter whether the packet within the IPComp-enabled flow is sent compressed or not. To indicate a packet is sent uncompressed, a new CPI value(TBD) is used. In this way, since all packets within the IPComp-enabled flow have

IPComp header, they will go through the same process path and be processed in order. For uncompressed packet, the Next Header in the IPComp Header is copied into the Next Header in the IP header, and the IPComp Header is removed.

5. IANA Considerations

This document require to add new CPI values in IKEv2 Notification IPCOMP Transform IDs (Value 16387) (<https://www.iana.org/assignments/ikev2-parameters/ikev2-parameters.xhtml>).

6. Security Considerations

The security requirements and mechanisms described in [RFC3173] also apply to this document.

This document does not introduce any new security considerations.

7. References

7.1. 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, <<https://www.rfc-editor.org/rfc/rfc2119>>.
- [RFC2407] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, DOI 10.17487/RFC2407, November 1998, <<https://www.rfc-editor.org/rfc/rfc2407>>.
- [RFC3051] Heath, J. and J. Border, "IP Payload Compression Using ITU-T V.44 Packet Method", RFC 3051, DOI 10.17487/RFC3051, January 2001, <<https://www.rfc-editor.org/rfc/rfc3051>>.
- [RFC3173] Shacham, A., Monsour, B., Pereira, R., and M. Thomas, "IP Payload Compression Protocol (IPComp)", RFC 3173, DOI 10.17487/RFC3173, September 2001, <<https://www.rfc-editor.org/rfc/rfc3173>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/rfc/rfc8174>>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/rfc/rfc8200>>.

7.2. Informative References

[CPI-IANA] "IPSEC IPCOMP Transform Identifiers", October 2022, <<https://www.iana.org/assignments/isakmp-registry/isakmp-registry.xhtml#isakmp-registry-11>>.

Appendix A. Contributors

The following individuals contributed significantly to this document:

Han Shi
Huawei Technologies
China
Email: shihang9@huawei.com

Authors' Addresses

Cheng Li
Huawei Technologies
China
Email: c.l@huawei.com

Guangming Zeng (editor)
Huawei Technologies
China
Email: zengguanming@huawei.com

Meng Zhang
Huawei Technologies
China
Email: zhangmeng6@huawei.com

Xiaobo Ding
Huawei Technologies
China
Email: mirroryuri.ding@huawei.com