

OPSAWG
Internet-Draft
Intended status: Standards Track
Expires: 5 June 2026

Y. Liu
T. Zhou
ZTE Corporation
2 December 2025

Export of Encapsulation Layer Information in IPFIX
draft-liu-opsawg-ipfix-muti-layer-01

Abstract

This document introduces new IPFIX IEs for encapsulation layer indication to help with the scenario when monitoring flow with multi-layer network encapsulations.

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 5 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	2
2. Terminology	3
3. Use Cases and Requirements	3
3.1. Use Cases	4
3.2. Requirements	5
4. IPFIX IEs for Encapsulation Layer Information	5
4.1. IPFIX Encapsulation Layer Information Option 1	6
4.1.1. encapLayerTop	6
4.1.2. encapLayer2	6
4.1.3. encapLayer3	7
4.2. IPFIX for Encapsulation Layer Information Option 2	7
5. Operational Considerations	10
5.1. Operational Considerations for Option 1	10
5.2. Operational Considerations for Option 2	11
6. Security Considerations	12
7. IANA Considerations	12
8. References	12
8.1. Normative References	12
8.2. Informative References	13
Authors' Addresses	14

1. Introduction

A packet may have multi-layer network encapsulations, and each layer may use the same or different network encapsulation headers. e.g, IP in IP encapsulation [RFC2003], which is an IP tunneling mechanism that encapsulates one IP packet in another IP packet, typical IP-in-IP scenario includes IPv4-in-IPv4, IPv6-in-IPv6, IPv4-in-IPv6 and IPv6-in-IPv4.

With the deployment of SRv6, the appearance of packets with IPv4-in-IPv6 or IPv6-in-IPv6 encapsulation becomes more and more common in the network. And there may be more than two network encapsulation layers in one packet as analyzed in section 3.1.

When monitoring a traffic flow with multiple encapsulations, e.g IP-in-IP, a typical use case is to answer the following questions:

- * Which tunnel are the packets steered into (e.g, identified by the outmost IP header) ?
- * What are the details of the inner packet (e.g, identified by the innermost IP header) ?

However, based on the existing IPFIX mechanisms, it is not easy to differentiate between IEs of different encapsulation layers.

This document aims to solve this problem by introducing new IPFIX IEs for encapsulation layer indication.

[Editor's Note] This version provides two alternative options to indicate the encapsulation layer information for discussion.

2. Terminology

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.

This document makes use of the terms defined in [RFC7011], [RFC8402] and [RFC8754].

The following terms are used as defined in [RFC7011]:

- * IPFIX
- * IPFIX Information Elements
- * Template
- * Collector
- * IPFIX Device

The following terms are used as defined in [RFC8402]:

- * Segment Routing (SR)
- * Segment List
- * SRv6
- * BSID


The following terms are used as defined in [RFC8754]:

- * SRH

3. Use Cases and Requirements

3.1. Use Cases

There may be more than two network encapsulation layers in one packet. As shown in the figure below.



CE1 sends IPv6 packets to CE2.

Packet leaving CE1: IPv6SA=CE1,DA=CE2>

PE1 performs SRv6 function H.Encaps with SRv6 Policy, and the corresponding SID list is <SID-R1,SID-T1,SID-PE2>, wherein SID-T1 is a BSID initiated on node T1.

Packet leaving PE1: IPv6<SA=PE1,DA=SID-R1><SID-R1,SID-T1,SID-PE2>IPv6<SA=CE1,DA=CE2>

When the packet arrives at T1, based on BSID SID-T1, T1 performs End.B6.Encaps [RFC8986] and encapsulate the third IPv6 header onto the packet with the corresponding SID list <SID-R2,SID-R3>.

Packet leaving T1:IPv6<SA=T1,DA=R2><SID-R2,SID-R3>IPv6<SA=PE1,DA=SID-PE2><SID-R1,SID-T1,SID-PE2>IPv6<SA=CE1,DA=CE2>

So the packet observed at node R2 has three IPv6 headers in it.

Based on different goals of the network monitor, the data collection scenario may include one of the following:

- * The network monitor wants to collect the information of the outmost SRH and the inner SRH of the same packet.
- * The network monitor only wants to collect the information of the outmost IPv6 header.
- * The network monitor only wants to collect the information of the innermost IPv6 header.
- * The network monitor wants to collect the SA of the outmost IPv6 header (i.e, the starting point of the tunnel) and the DA of the innermost IPv6 header(i.e, the final destination of the packet)

3.2. Requirements

Based on the scenarios described in section 3.1, the information collection requirements in IPFIX for multi-layer encapsulated packets include:

- * Req-a: Collecting the same fields from both the outer and inner packet headers at the same time.
- * Req-b: Collecting only the fields from the inner packet header.
- * Req-c: Collecting only the fields from the outer packet header
- * Req-d: Collecting different fields from the outer and inner packet headers at the same time.

Req-a can be fulfilled leveraging the existing mechanism. As described in [RFC7011], if the same element appears multiple times in an IPFIX template, it should be processed in order. For example, when exporting the two source IP addresses of an IPv6-in-IPv6 packet, the first sourceIPv6Address Information Element occurrence should be the IPv6 address of the outer header, while the second occurrence should be the address of the inner header.

However, Req-b, Req-b and Req-d can not be meet using standard method by far. When receiving a IPFIX message with a certain IE(e.g, sourceIPv6Address) from the IPFIX Device, the collector is not able to tell which layer this IE belongs to for traffic flow with multi-layer encapsulations.

4. IPFIX IEs for Encapsulation Layer Information

In the following subsections, two options are provided for encapsulation layer indication:

- * Option 1: New IEs are defined as the layer separators. The meaning of the IEs of Header Fields (in other words, which encapsulation layer these IEs belongs to), depends on the appearance and the content of the "layer separator" IEs.
- * Option 2: A new IE using the data type "subTemplateMultiList" [RFC6313] is defined to carry the multiple encapsulation packet header fields as the structured data. This method avoids the dependencies between IEs.

4.1. IPFIX Encapsulation Layer Information Option 1

This section defines several Encapsulation Layer IEs for network encapsulation layer indication. When there is no need to differentiate between these IEs in this document, they will be collectively referred to as "encapsulation layer IE".

4.1.1. encapLayerTop

A new IE "encapLayerTop" is defined in this section to indicate which IEs in the IPFIX messages belongs to the outmost network encapsulation layer.

Name: encapLayerTop

ElementID: TBD1

Description: A 16-bit identifier indicating that the IEs follows immediately after it till the next Encapsulation Layer IE belong to the outmost network encapsulation layer (e.g, from the outmost Ethernet header to the first IP header). If there's not any other Encapsulation Layer IE exists in the Template, it means that all the IEs following encapLayerTop belong to the outmost network encapsulation layer. This IE has a fixed value of 0xFFFF.

Abstract Data Type: unsigned16

Data Type Semantics: identifier

Reference: This document.

4.1.2. encapLayer2

A new IE "encapLayer2" is defined in this section to indicate which IEs in the IPFIX messages belongs to the second network encapsulation layer.

Name: encapLayer2

ElementID: TBD2

Description: A 16-bit identifier indicating that the IEs follows immediately after it till the next Encapsulation Layer IE belong to the second network encapsulation layer. If there's not any other Encapsulation Layer IE exists in the Template, it means that all the IEs following encapLayer2 belong to the second network encapsulation layer. This IE has a fixed value of 0xFFFF.

Abstract Data Type: unsigned16

Data Type Semantics: identifier

Reference: This document.

4.1.3. encapLayer3

A new IE "encapLayer3" is defined in this section to indicate which IEs in the IPFIX messages belongs to the third network encapsulation layer.

Name: encapLayer3

ElementID: TBD2

Description: A 16-bit identifier indicating that the IEs follows immediately after it till the next Encapsulation Layer IE belong to the third network encapsulation layer. If there's not any other Encapsulation Layer IE exists in the Template, it means that all the IEs following encapLayer3 belong to the third network encapsulation layer. This IE has a fixed value of 0xFFFF.

Abstract Data Type: unsigned16

Data Type Semantics: identifier

Reference: This document.

4.2. IPFIX for Encapsulation Layer Information Option 2

A new IE "multiLayerEcapFields" is defined in this section as shown below. This IE supports to carry the header fields for packets with multiple layers of encapsulation.

Name: multiLayerEcapFields

ElementID: TBA1

Description: * An IPFIX Information Element that denotes that the header fields of different encapsulation layers will be exported. Each top-level element in a subTemplateMultiList Information Element carries a Template ID, Length, and zero or more Data Records corresponding to the Template ID. The "ordered" structured data type semantic [RFC6313] is used. The Length field of this subTemplateMultiList is encoded in three bytes even though it may be less than 255 octets.

- * The template IDs from top to bottom carried in the IE correspond to encapsulation layers of the packet flow, starting from the outermost layer. For example, the data record of the template whose template ID is carried in the first list element belongs to the outmost network encapsulation layer, the data record of the template whose template ID is carried in the second list element belongs to the second network encapsulation layer, and the data record of the template whose template ID is carried in the third list element belongs to the third network encapsulation layer.
- * If not all of the encapsulation layers are required to be exported, e.g., only the information of layer 1 and layer 3 needs to be exported, the template ID in the data record and the corresponding length field **MUST** be set to 0 to indicate the absence of the information of this layer.

Abstract Data Type: subTemplateList

Data Type Semantics: list

Reference: This document.

An example of the encoding of the IE "multiLayerEcapFields" is shown below for a traffic flow with IPv6-in-IPv6-in-IPv6 encapsulation, where the destination address of the outmost IPv6 header and the source address of the innermost IPv6 header are required to be exported, while the information of the second IPv6 header is not needed.




Figure 1: Encoding multiLayerEcapFields, Template for Flow Record


```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Set ID = 3           |           Length = 12           |
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Template ID = 301    |           Field Count = 1       |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 | IE=destinationIPv6Address(28) |           Field Length = 16   |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 2: Template for the outmost IPv6 header

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Set ID = 4           |           Length = 12           |
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Template ID = 302    |           Field Count = 1       |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 | IE = sourceIPv6Address(27)  |           Field Length = 16   |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 3: Template for the third IPv6 header

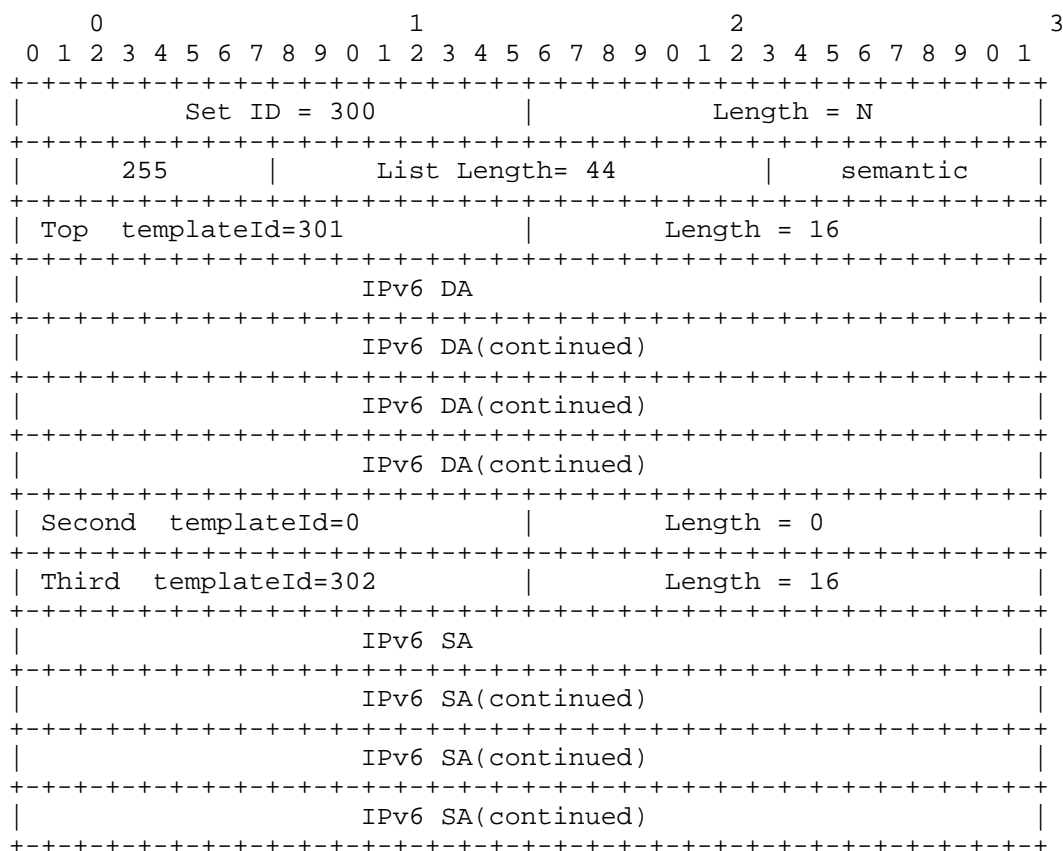


Figure 4: Encoding multiple-layer headers, Data Set

5. Operational Considerations

It should be noticed that the layer information on different exporters may be different, since the networking context might be different from each exporter and each exporter only knows the layer information locally.

5.1. Operational Considerations for Option 1

To generate Flow Records with IEs for encapsulation layer included, the metering process SHOULD recognize the encapsulation layer of the corresponding fields in the packet. This is mainly based on local implementation and the details are out of the scope of this document.

The order of the IEs in the Template Records MUST be guaranteed when using encapLayerTop/encapLayer2/encapLayer3, that is, the IEs belonging to each encapsulation layer MUST immediately follow the corresponding encapsulation layer IE.

Each encapsulation layer IE SHALL NOT appear more than once in a Template. If there's more than one encapsulation layer IE of the same type in the Template, the Collecting Process MUST ignore the Template and the Collecting Process SHOULD log the error.

As in [RFC5012], the Information Elements are derived from fields of packets or from packet treatment. For IEs that are not related with header fields, whether they are covered by scope of the encapsulation layer IE, they SHOULD be processed following the existing specifications.

For IEs of Header Fields that are not in the scope of encapsulation layer IE, e.g, there're IEs of Header Fields in the Template before the appearance of Encapsulation Layer IEs, they SHOULD be processed properly based on the default behavior of the Collector, how the Collector would process them is out of the scope of this document.

Currently only three IEs are defined to meet the traffic monitoring requirements for packets of no more than three layers. If in the future , there're packets with four or five or more encapsulation layers in the network and the packet header information of each layer is required to be exported, more new IEs such as encapLayer4/encapLayer5 MAY be needed

5.2. Operational Considerations for Option 2

Beside the templateId 0, if the templateId used by the multiLayerEcapFields IE in the data record doesn't exit, the collector SHOULD ignore the encapsulation layer this templateId represents and log an error, and it is RECOMMENDED to process the information of other layers normally.

For IEs of Header Fields that are not included the multiLayerEcapFields IE, they SHOULD be processed properly based on the default behavior of the Collector, how the Collector would process them is out of the scope of this document.

As in [RFC5012], the Information Elements are derived from fields of packets or from packet treatment. For IEs that are not related with header fields, whether they are included in the encapsulation layer IE, they SHOULD be processed following the existing specifications.

6. Security Considerations

There are no additional security considerations regarding allocation of these new IPFIX IEs compared to [RFC7012].

7. IANA Considerations

For option 1, this document requests IANA to create new IEs under the "IPFIX Information Elements" registry [RFC7012] available at [IANA-IPFIX].

Element ID	Name	Reference
TBD1	encapLayerTop	This document
TBD2	encapLayer2	This document
TBD3	encapLayer3	This document

For option 2, this document requests IANA to create new IEs under the "IPFIX Information Elements" registry [RFC7012] available at [IANA-IPFIX].

Element ID	Name	Reference
TBA1	multiLayerEcapFields	This document

8. References

8.1. Normative References

[IANA-IPFIX]

IANA, "IP Flow Information Export (IPFIX) Entities",
<<https://www.iana.org/assignments/ipfix>>.

[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/info/rfc2119>>.

- [RFC6313] Claise, B., Dhandapani, G., Aitken, P., and S. Yates, "Export of Structured Data in IP Flow Information Export (IPFIX)", RFC 6313, DOI 10.17487/RFC6313, July 2011, <<https://www.rfc-editor.org/info/rfc6313>>.
- [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <<https://www.rfc-editor.org/info/rfc7011>>.
- [RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, DOI 10.17487/RFC7012, September 2013, <<https://www.rfc-editor.org/info/rfc7012>>.
- [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/info/rfc8174>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

8.2. Informative References

- [RFC2003] Perkins, C., "IP Encapsulation within IP", RFC 2003, DOI 10.17487/RFC2003, October 1996, <<https://www.rfc-editor.org/info/rfc2003>>.
- [RFC5012] Schulzrinne, H. and R. Marshall, Ed., "Requirements for Emergency Context Resolution with Internet Technologies", RFC 5012, DOI 10.17487/RFC5012, January 2008, <<https://www.rfc-editor.org/info/rfc5012>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

Authors' Addresses

Yao Liu
ZTE Corporation
China
Email: liu.yao71@zte.com.cn

Taoran Zhou
ZTE Corporation
China
Email: zhou.taoran@zte.com.cn