

OPSAWG  
Internet-Draft  
Intended status: Standards Track  
Expires: 1 June 2026

Y. Liu  
L. Zhao  
ZTE  
28 November 2025

Export of BGP VPN Information in IPFIX  
draft-liu-opsawg-ipfix-bgp-vpn-00

## Abstract

This document introduces new IP Flow Information Export (IPFIX) information elements to carry the egress PE information in IPFIX.

## 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 1 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. New IPFIX IEs for VPN Egress PE Information . . . . . | 4  |
| 3.1. BGP VPN Next Hop Information . . . . .              | 4  |
| 3.1.1. bgpVpnNextHopIPv4Address . . . . .                | 4  |
| 3.1.2. bgpVpnNextHopIPv6Address . . . . .                | 5  |
| 3.2. SRv6 Service SID Locator in IPFIX . . . . .         | 5  |
| 3.2.1. srv6ServiceSidLocator . . . . .                   | 6  |
| 3.2.2. srv6ServiceSidLocatorLength . . . . .             | 6  |
| 4. Operational Considerations . . . . .                  | 7  |
| 5. Security Considerations . . . . .                     | 7  |
| 6. IANA Considerations . . . . .                         | 7  |
| 7. References . . . . .                                  | 8  |
| 7.1. Normative References . . . . .                      | 8  |
| 7.2. Informative References . . . . .                    | 9  |
| Authors' Addresses . . . . .                             | 10 |

## 1. Introduction

BGP/MPLS VPN, as described in [RFC4364], is a method that uses BGP to exchange the routes of a particular VPN among the PE routers that are attached to that VPN. And each route within a VPN is assigned an MPLS label.

Typical MPLS VPN scenarios include:

- \* MPLS VPN over MPLS traffic engineering (MPLS-TE) tunnel: The MPLS-TE tunnel can be built based on RSVP-TE LSP [RFC5824] or SR-MPLS Policy.
- \* MPLS VPN with MPLS best effort tunnel: A single MPLS label/SR-MPLS SID for the FEC on the egress PE is used to tunnel the VPN traffic over the backbone.

For SRv6 VPN services, [RFC9252] defines procedures and messages for SRv6-based BGP services, including L3VPN, EVPN, and Internet services. SRv6 Service SID refers to an SRv6 SID associated with one of the service-specific SRv6 Endpoint Behaviors on the advertising PE router.

As in [RFC9252], typical SRv6 VPN scenario includes:

- \* SRv6 service with SRv6-TE connectivity: To provide SRv6 service in conjunction with an underlay Service Level Agreement (SLA) from the ingress PE to the egress PE, the egress PE colors the overlay service route with a Color Extended Community [RFC9012] for

steering flows for those routes. The ingress PE encapsulates the payload packet in an outer IPv6 header with the SR Policy segment list associated with the related SLA along with the SRv6 Service SID associated with the route using the Segment Routing Header (SRH) [RFC8754].

- \* SRv6 service with best-effort(SRv6-BE) connectivity: The egress PE signals an SRv6 Service SID with the BGP overlay service route. The ingress PE encapsulates the payload in an outer IPv6 header where the destination address is the SRv6 Service SID provided by the egress PE. The underlay between the PEs only needs to support plain IPv6 forwarding.

When monitoring traffic flows on the ingress PE in a network with BGP VPN deployed, the network monitor may want to know the following information:

- \* Which egress PE is the flow forwarded to ?
- \* How is the traffic transmitted through the network ?

This document introduces new IP Flow Information Export (IPFIX) information elements to carry the egress PE information in IPFIX.

## 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 [RFC9252].

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

- \* IPFIX
- \* IPFIX Information Elements
- \* Metering Process
- \* Template Record
- \* Data Record
- \* Collector

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

- \* Segment Routing (SR)
- \* Segment List
- \* SRv6
- \* SR-MPLS
- \* Segment Identifier (SID)

The following terms are used as defined in [RFC9252]:

- \* SRv6 Service SID

### 3. New IPFIX IEs for VPN Egress PE Information

The following subsections defines different types of IEs to fulfill the requirement to obtain the egress PE information via IPFIX.

#### 3.1. BGP VPN Next Hop Information

Two new IEs are defined in this section to identify the next hop address of the BGP VPN route. One for IPv4 address and the other for IPv6 address. The BGP next hop address is an address of the egress PE router as in [RFC4364].

##### 3.1.1. bgpVpnNextHopIPv4Address

Name: bgpVpnNextHopIPv4Address

ElementID: TBD1

Description: The 32-bit IPv4 address on the egress PE which is used as the next hop address of the BGP VPN route.

Abstract Data Type: default

Data Type Semantics: ipv4Address

Additional Information: Specified in [RFC4364].

Reference: This document.

### 3.1.2. bgpVpnNextHopIPv6Address

Name: bgpVpnNextHopIPv6Address

ElementID: TBD2

Description: The 128-bit IPv6 address on the egress PE which is used as the next hop address of the BGP VPN route.

Abstract Data Type: default

Data Type Semantics: ipv6Address

Additional Information: See [RFC4659] for more information about the IPv6 Next Hop Network Address.

Reference: This document.

### 3.2. SRv6 Service SID Locator in IPFIX

In the case of SRv6 VPN, another choice to be aware of the egress PE information is to export the locator information of the SRv6 service SID, since generally the SRv6 locators are well planned in the network, and different PEs are usually assigned with different locators.

[RFC9487] defines IE "srhSegmentIPv6" and IE "srhSegmentIPv6LocatorLength", and it enables the calculation of the SRv6 Locator when the two IEs are used together. However, the requirement to export the locator of the SRv6 service SID can not be fulfilled using "srhSegmentIPv6" and "srhSegmentIPv6LocatorLength" due to the following reasons:

- \* In the SRv6-TE scenario, the SRv6 service SID would be encapsulated in the SRH as the last segment(i.e, Segment List[0]) of the segment list in SRH. Although "srhSegmentIPv6" is the 128-bit IPv6 address that represents an SRv6 segment, there's no mechanism yet to solely export Segment List[0](or any other segment besides the active segment) in the SRH.

- \* In the SRv6-BE scenario, the SRv6 service SID is encapsulated as the destination address of the IPv6 header by the ingress PE. Theoretically, the IE "destinationIPv6Address" and "destinationIPv6PrefixLength" defined in [RFC7012] can be used to calculate the the IPv6 prefix length of the SRv6 service SID. But if this method is used, the network analyzer needs to know exactly which flows are VPN flows using SRv6-BE forwarding to distinguish SRv6 Service SID from the normal IPv6 address carried in the IPv6 destination address field.

To export locator of the SRv6 Service SID which is advertised via BGP VPN routes, the following IEs are defined, and this method is applicable for both SRv6-TE and SRv6-BE scenario.

### 3.2.1. srv6ServiceSidLocator

Name: srv6ServiceSidLocator

ElementID: TBD3

Description: The Locator of the SRv6 Service SID signaled by the egress PE via BGP.

Abstract Data Type: default

Data Type Semantics: ipv6Address

Additional Information: See [RFC9252] for more information about the SRv6 service SID. See Section 3.1 of [RFC8986] for more details about the SID format.

Reference: This document.

### 3.2.2. srv6ServiceSidLocatorLength

Name: srv6ServiceSidLocatorLength

ElementID: TBD4

Description: The length of the SRv6 Locator of the SRv6 service SID specified as the number of significant bits. Together with srv6ServiceSid, it enables the calculation of SRv6 Locator of the SRv6 service SID.

Abstract Data Type: default

Data Type Semantics: default

Additional Information: See Section 3.1 of [RFC8986] for more details about the SID format.

Reference: This document.

#### 4. Operational Considerations

The IE `bgpNextHopIPv4Address(18)` and `bgpNextHopIPv6Address(63)` define the IPv4/IPv6 address of the next (adjacent) BGP hop. If BGP VPN route is the only BGP route deployed on the PE, IE 18 and IE 63 MAY be used to indicate the next hop address of the BGP VPN route. However, when there're many types of BGP route used in the network(e.g., BGP VPN [RFC4364] is used together with BGP-LU[RFC8277]), it is not clear which type of the BGP route the next BGP hop carried in IE 18 or IE 63 belongs to. In this case, using `bgpVpnNextHopIPv4Address` and `bgpVpnNextHopIPv6Address` defined in this document to carry the next hop address of the BGP VPN route is more appropriate.

In the multi-as backbones, if inter-AS option A or option B with BGP next-hop changed are used as described in Section 10 of [RFC4364], the address of the egress PE can't be obtained via "`bgpVpnNextHopIPv4Address`" or "`bgpVpnNextHopIPv6Address`" since the next hop address of the BGP VPN route received by the ingress PE is not the address of the egress PE.

#### 5. Security Considerations

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

Other security considerations for BGP/MPLS VPN in [RFC4364] and for BGP Overlay Services Based on SRv6 in [RFC9252] apply to this document.

#### 6. IANA Considerations

This document requests IANA to create new IEs under the "IPFIX Information Elements" registry [RFC7012] available at [IANA-IPFIX].

| Element ID | Name                        | Reference     |
|------------|-----------------------------|---------------|
| TBD1       | bgpVpnNextHopIPv4Address    | Section 3.1.1 |
| TBD2       | bgpVpnNextHopIPv6Address    | Section 3.1.2 |
| TBD3       | srv6ServiceSidLocator       | Section 3.2.1 |
| TBD4       | srv6ServiceSidLocatorLength | Section 3.2.2 |

Table 1: IPFIX Information Elements Registry

## 7. References

### 7.1. 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>>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.

[RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, DOI 10.17487/RFC4659, September 2006, <<https://www.rfc-editor.org/info/rfc4659>>.

[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>>.
- [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>>.
- [RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <<https://www.rfc-editor.org/info/rfc9252>>.

## 7.2. Informative References

- [RFC5824] Kumaki, K., Ed., Zhang, R., and Y. Kamite, "Requirements for Supporting Customer Resource ReSerVation Protocol (RSVP) and RSVP Traffic Engineering (RSVP-TE) over a BGP/MPLS IP-VPN", RFC 5824, DOI 10.17487/RFC5824, April 2010, <<https://www.rfc-editor.org/info/rfc5824>>.
- [RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/info/rfc8277>>.
- [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>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9487] Graf, T., Claise, B., and P. Francois, "Export of Segment Routing over IPv6 Information in IP Flow Information Export (IPFIX)", RFC 9487, DOI 10.17487/RFC9487, November 2023, <<https://www.rfc-editor.org/info/rfc9487>>.

Authors' Addresses

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

Liman Zhao  
ZTE  
Email: zhao.liman@zte.com.cn