

Internet Engineering Task Force (IETF)
Request for Comments: 9160
Category: Informational
ISSN: 2070-1721

T. Graf
Swisscom
December 2021

Export of MPLS Segment Routing Label Type Information in IP Flow Information Export (IPFIX)

Abstract

This document introduces new IP Flow Information Export (IPFIX) code points to identify which traffic is being forwarded based on which MPLS control plane protocol is used within a Segment Routing domain. In particular, this document defines five code points for the IPFIX `mplsTopLabelType` Information Element for Path Computation Element (PCE), IS-IS, OSPFv2, OSPFv3, and BGP MPLS Segment Routing extensions.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9160>.

Copyright Notice

Copyright (c) 2021 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.	MPLS Segment Routing Top Label Type
3.	IANA Considerations
4.	Operational Considerations
5.	Security Considerations
6.	References
6.1.	Normative References
6.2.	Informative References
	Acknowledgements
	Author's Address

1. Introduction

Four routing protocol extensions -- OSPFv2 Extensions [RFC8665], OSPFv3 Extensions [RFC8666], IS-IS Extensions [RFC8667], and BGP Prefix Segment Identifiers (Prefix-SIDs) [RFC8669] -- and one Path Computation Element Communication Protocol (PCEP) Extension [RFC8664] have been defined to be able to propagate Segment Routing (SR) labels for the MPLS data plane [RFC8660].

Also, [SR-Traffic-Accounting] describes how IP Flow Information Export (IPFIX) [RFC7012] can be leveraged in dimensional data modeling to account for traffic to MPLS SR label dimensions within a Segment Routing domain.

In [RFC7012], the Information Element (IE) `mplsTopLabelType(46)` identifies which MPLS control plane protocol allocated the top-of-stack label in the MPLS label stack. Per Section 7.2 of [RFC7012], the "IPFIX MPLS label type (Value 46)" subregistry [IANA-IPFIX] was created, where new MPLS label type entries should be added. This document defines new code points to address typical use cases that are discussed in Section 2.

2. MPLS Segment Routing Top Label Type

By introducing five new code points to the IPFIX IE `mplsTopLabelType(46)` for Path Computation Element (PCE), IS-IS, OSPFv2, OSPFv3, and BGP Prefix-SIDs, it is possible to identify which traffic is being forwarded based upon which MPLS SR control plane protocol is in use.

A typical use case is to monitor MPLS control plane migrations from LDP to IS-IS or OSPF Segment Routing. Such a migration can be done node by node as described in Appendix A of [RFC8661].

Another use case is to monitor MPLS control plane migrations from dynamic BGP labels [RFC8277] to BGP Prefix-SIDs [RFC8669]. For example, the motivation for, and benefits of, such a migration in large-scale data centers are described in [RFC8670].

Both use cases can be verified by using `mplsTopLabelType(46)`, `mplsTopLabelIPv4Address(47)`, `mplsTopLabelIPv6Address(140)`, `mplsTopLabelStackSection(70)`, and `forwardingStatus(89)` IEs to infer

- * how many packets are forwarded or dropped
- * if packets are dropped, for which reasons, and
- * the MPLS provider edge loopback address and label protocol

By looking at the MPLS label value itself, it is not always clear to which label protocol it belongs. This is because they may share the same label allocation range. This is, for example, the case for IGP-Adjacency SIDs, LDP, and dynamic BGP labels.

3. IANA Considerations

IANA has allocated the following code points in the "IPFIX MPLS label type (Value 46)" subregistry within the "IPFIX Information Elements" registry [RFC7012]. See [IANA-IPFIX].

Value	Description	Reference
6	Path Computation Element	RFC 9160, RFC 8664

7	OSPFv2 Segment Routing	RFC 9160, RFC 8665	
+-----+	+-----+	+-----+	+-----+
8	OSPFv3 Segment Routing	RFC 9160, RFC 8666	
+-----+	+-----+	+-----+	+-----+
9	IS-IS Segment Routing	RFC 9160, RFC 8667	
+-----+	+-----+	+-----+	+-----+
10	BGP Segment Routing Prefix-SID	RFC 9160, RFC 8669	
+-----+	+-----+	+-----+	+-----+

Table 1: Updates to "IPFIX MPLS label type (Value 46)"
Subregistry

References to RFCs 4364, 4271, and 5036 have been added to the "Reference" column in the "IPFIX MPLS label type (Value 46)" subregistry [IANA-IPFIX] for code points 3, 4, and 5, respectively. Previously, these references appeared in the "Additional Information" column for `mplsTopLabelType(46)` in the "IPFIX Information Elements" registry [IANA-IPFIX].

4. Operational Considerations

In the IE `mplsTopLabelType(46)`, BGP code point 4 refers to the label value in the `MP_REACH_NLRI` path attribute described in Section 2 of [RFC8277], while BGP Segment Routing Prefix-SID code point 10 corresponds to the label index value in the Label-Index TLV described in Section 3.1 of [RFC8669]. These values are thus used for those distinct purposes.

5. Security Considerations

There exist no significant extra security considerations regarding the allocation of these new IPFIX IEs as compared to [RFC7012].

6. References

6.1. Normative References

[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>>.

6.2. Informative References

[IANA-IPFIX]
IANA, "IPFIX MPLS label type (Value 46)", <<https://www.iana.org/assignments/ipfix/>>.

[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>>.

[RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.

[RFC8661] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., and S. Litkowski, "Segment Routing MPLS Interworking with LDP", RFC 8661, DOI 10.17487/RFC8661, December 2019, <<https://www.rfc-editor.org/info/rfc8661>>.

[RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664,

DOI 10.17487/RFC8664, December 2019,
<<https://www.rfc-editor.org/info/rfc8664>>.

[RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.

[RFC8666] Psenak, P., Ed. and S. Previdi, Ed., "OSPFv3 Extensions for Segment Routing", RFC 8666, DOI 10.17487/RFC8666, December 2019, <<https://www.rfc-editor.org/info/rfc8666>>.

[RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.

[RFC8669] Previdi, S., Filsfils, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <<https://www.rfc-editor.org/info/rfc8669>>.

[RFC8670] Filsfils, C., Ed., Previdi, S., Dawra, G., Aries, E., and P. Lapukhov, "BGP Prefix Segment in Large-Scale Data Centers", RFC 8670, DOI 10.17487/RFC8670, December 2019, <<https://www.rfc-editor.org/info/rfc8670>>.

[SR-Traffic-Accounting]

Ali, Z., Filsfils, C., Talaulikar, K., Sivabalan, S., Horneffer, M., Raszuk, R., Litkowski, S., Voyer, D., Morton, R., and G. Dawra, "Traffic Accounting in Segment Routing Networks", Work in Progress, Internet-Draft, draft-ali-spring-sr-traffic-accounting-06, 13 November 2021, <<https://datatracker.ietf.org/doc/html/draft-ali-spring-sr-traffic-accounting-06>>.

Acknowledgements

I would like to thank the IE doctors, Paul Aitken and Andrew Feren, as well as Benot Claise, Loa Andersson, Tianran Zhou, Pierre Franois, Bruno Decraene, Paolo Lucente, Hannes Gredler, Ketan Talaulikar, Sabrina Tanamal, Erik Auerswald, Sergey Fomin, Mohamed Boucadair, Tom Petch, Qin Wu, and Matthias Arnold for their review and valuable comments. Many thanks also to Robert Wilton for the AD review. Thanks to Alvaro Retana, ric Vyncke, and Benjamin Kaduk for the IESG review.

Author's Address

Thomas Graf
Swisscom
Binzring 17
CH-8045 Zrich
Switzerland

Email: thomas.graf@swisscom.com