

Workgroup: Network Working Group
Internet-Draft:
draft-ietf-opsawg-ipfix-mpls-sr-label-type-11
Published: 18 September 2021
Intended Status: Informational
Expires: 22 March 2022
Authors: T. Graf

Swisscom

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 used within a Segment Routing domain. In particular, this document defines five code points for the IPFIX mplsTopLabelType Information Element for PCE, IS-IS, OSPFv2, OSPFv3, and BGP MPLS Segment Routing extensions.

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 22 March 2022.

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 Simplified BSD License text as described in

Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#)
- [2. MPLS Segment Routing Top Label Type](#)
- [3. IANA Considerations](#)
- [4. Operational Considerations](#)
- [5. Security Considerations](#)
- [6. Acknowledgements](#)
- [7. References](#)
 - [7.1. Normative References](#)
 - [7.2. Informative References](#)
- [Author's Address](#)

1. Introduction

Four routing protocol extensions, [OSPFv2 Extensions](#) [[RFC8665](#)], [OSPFv3 Extensions](#) [[RFC8666](#)], [IS-IS Extensions](#) [[RFC8667](#)], [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, [[I-D.ali-spring-sr-traffic-accounting](#)] describes how IP Flow Information Export [[RFC7012](#)] can be leveraged in dimensional data modelling to account 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. Section 7.2 of [[RFC7012](#)] creates the "[IPFIX MPLS label type \(Value 46\)](#)" subregistry [[IANA-IPFIX](#)] where MPLS label type 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 PCE, IS-IS, OSPFv2, OSPFv3 and BGP Prefix-SID, 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 and benefits for 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 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 as 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

This document requests IANA to allocate the following code points in the existing subregistry "IPFIX MPLS label type (Value 46)" under the "IPFIX Information Elements" registry [[RFC7012](#)] available at [[IANA-IPFIX](#)].

Value	Description	Reference
TBD1	Path Computation Element	[RFC-to-be], RFC8664
TBD2	OSPFv2 Segment Routing	[RFC-to-be], RFC8665
TBD3	OSPFv3 Segment Routing	[RFC-to-be], RFC8666
TBD4	IS-IS Segment Routing	[RFC-to-be], RFC8667
TBD5	BGP Segment Routing Prefix-SID	[RFC-to-be], RFC8669

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

Note to the RFC-Editor:

- *Please replace TBD1 - TBD5 with the values allocated by IANA

*Please replace the [RFC-to-be] with the RFC number assigned to this document

Note IANA:

*Suggest to move the existing RFC references in the additional information column of IE mplsTopLabelType(46) to reference column for codepoint 3, 4 and 5.

4. Operational Considerations

In the IE mplsTopLabelType(46), the BGP code point 4 refers to the label value in MP_REACH_NLRI path attribute described in Section 2 of [RFC8277], while the BGP Segment Routing Prefix-SID code point TBD5 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 exists no significant extra security considerations regarding the allocation of these new IPFIX IEs compared to [RFC7012].

6. Acknowledgements

I would like to thank the IE doctors, Paul Aitken and Andrew Feren, as well Benoit Claise, Loa Andersson, Tianran Zhou, Pierre Francois, Bruno Decreane, 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, Eric Vyncke and Benjamin Kaduk for the IESG review.

7. References

7.1. 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>>.

7.2. Informative References

[I-D.ali-spring-sr-traffic-accounting]

Filsfils, C., Talaulikar, K., Sivabalan, S., Horneffer, M., Raszuk, R., Litkowski, S., Voyer, D., and R. Morton, "Traffic Accounting in Segment Routing Networks", Work in Progress, Internet-Draft, draft-ali-spring-sr-traffic-

accounting-05, 12 April 2021, <<https://www.ietf.org/archive/id/draft-ali-spring-sr-traffic-accounting-05.txt>>.

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

[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>>.

Author's Address

Thomas Graf
Swisscom
Binzring 17
CH-8045 Zurich
Switzerland

Email: thomas.graf@swisscom.com