
Stream: Internet Engineering Task Force (IETF)
RFC: [9714](#)
Category: Standards Track
Published: January 2025
ISSN: 2070-1721
Authors: W. Cheng, Ed. X. Min, Ed. T. Zhou J. Dai Y. Peleg
China Mobile ZTE Corp. Huawei FiberHome Broadcom

RFC 9714

Encapsulation for MPLS Performance Measurement with the Alternate-Marking Method

Abstract

This document defines the encapsulation for MPLS performance measurement with the Alternate-Marking Method, which performs flow-based packet loss, delay, and jitter measurements on MPLS traffic.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in 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/rfc9714>.

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. Conventions Used in This Document	3
2.1. Abbreviations	3
2.2. Requirements Language	4
3. Flow-Based PM Encapsulation in MPLS	4
3.1. Examples for Applying Flow-ID Label in a Label Stack	6
3.1.1. Layout of the Flow-ID Label when Applied to MPLS Transport	6
3.1.2. Layout of the Flow-ID Label when Applied to MPLS Service	6
3.1.3. Layout of the Flow-ID Label when Applied to both MPLS Transport and MPLS Service	7
4. Procedures of Encapsulation, Look-Up, and Decapsulation	8
5. Procedures of Flow-ID Allocation	9
6. FLC and FRLD Considerations	10
7. Equal-Cost Multipath Considerations	10
8. Security Considerations	10
9. IANA Considerations	11
10. References	11
10.1. Normative References	11
10.2. Informative References	12
Acknowledgements	13
Contributors	13
Authors' Addresses	13

1. Introduction

[RFC9341] describes a performance measurement method, which can be used to measure packet loss, delay, and jitter on data traffic. Since this method is based on marking consecutive batches of packets, it is referred to as the Alternate-Marking Method. [RFC8372] outlines key considerations for developing a solution for MPLS flow identification, intended for use in performance monitoring of MPLS flows.

This document defines the encapsulation for MPLS performance measurement with the Alternate-Marking Method, which performs flow-based packet loss, delay, and jitter measurements on the MPLS traffic. The encapsulation defined in this document supports performance monitoring at the intermediate nodes and MPLS flow identification at both transport and service layers.

Note that in parallel to the work of this document, there is ongoing work on MPLS Network Actions (MNA) [RFC9613]. The MPLS performance measurement with the Alternate-Marking Method can also be achieved by MNA encapsulation. In addition, MNA will provide a broader use-case applicability. That means the MNA encapsulation is expected to provide a more advanced solution. Once published as an RFC, it is agreed that this document will be made Historic.

2. Conventions Used in This Document

2.1. Abbreviations

ACL: Access Control List

BoS: Bottom of Stack

cSPL: Composite Special Purpose Label, the combination of the Extension Label (value 15) and an Extended Special Purpose Label

DSCP: Differentiated Services Code Point

ELC: Entropy Label Capability

ERLD: Entropy Readable Label Depth

eSPL: Extended Special Purpose Label, a special-purpose label that is placed in the label stack after the Extension Label (value 15)

FL: Flow-ID Label

FLC: Flow-ID Label Capability

FLI: Flow-ID Label Indicator

FRLD: Flow-ID Readable Label Depth

IPFIX: IP Flow Information Export [[RFC7011](#)]

LSP: Label Switched Path

LSR: Label Switching Router

MPLS: Multi-Protocol Label Switching

NMS: Network Management System

PHP: Penultimate Hop Popping

PM: Performance Measurement

PW: PseudoWire

SFL: Synonymous Flow Label

SID: Segment ID

SR: Segment Routing

TC: Traffic Class

TTL: Time to Live

VC: Virtual Channel

VPN: Virtual Private Network

XL: Extension Label

2.2. 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. Flow-Based PM Encapsulation in MPLS

This document defines the Flow-based MPLS performance measurement encapsulation with the Alternate-Marking Method, as shown in [Figure 1](#).

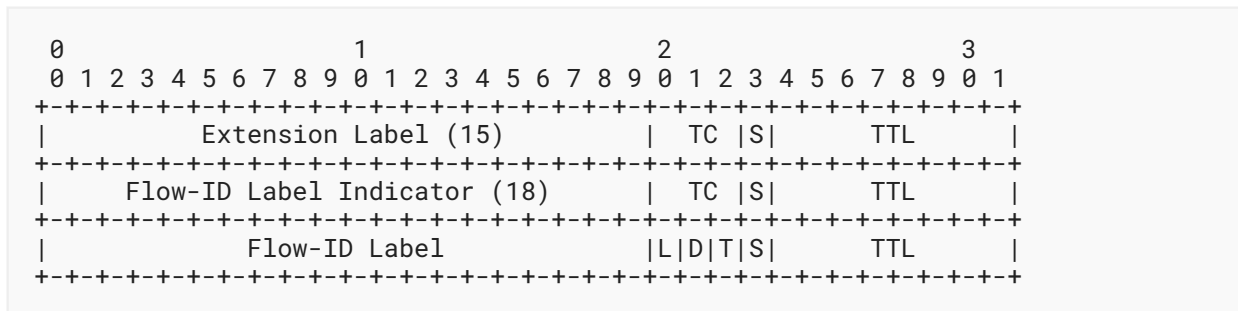


Figure 1: Flow-based PM Encapsulation in MPLS

The Flow-ID Label Indicator (FLI) is an Extended Special Purpose Label (eSPL), which is combined with the Extension Label (XL, value 15) to form a Composite Special Purpose Label (cSPL), as defined in [RFC9017]. The FLI is defined in this document as value 18.

The Traffic Class (TC) and Time To Live (TTL) fields of the XL and FLI **MUST** use the same values of the label immediately preceding the XL. The Bottom of the Stack (BoS) bit [RFC3032] for the XL and FLI **MUST** be zero. If any XL or FLI processed by a node has the BoS bit set, the node **MUST** discard the packet and **MAY** log an error.

The Flow-ID Label (FL) is used as an MPLS flow identification [RFC8372]. Its value **MUST** be unique within the administrative domain. The FL values **MAY** be allocated by an external NMS or controller based on the measurement object instances (such as LSP or PW). There is a one-to-one mapping between a Flow-ID and a flow. The specific method on how to allocate the FL values is described in Section 5.

The FL, preceded by a cSPL, can be placed either at the bottom or in the middle, but not at the top, of the MPLS label stack, and it **MAY** appear multiple times within a label stack. Section 3.1 of this document provides several examples to illustrate the application of FL in a label stack. The TTL for the FL **MUST** be zero to ensure that it is not used inadvertently for forwarding. The BoS bit for the FL depends on whether the FL is placed at the bottom of the MPLS label stack, i.e., the BoS bit for the FL is set only when the FL is placed at the bottom of the MPLS label stack.

Besides the flow identification, a color-marking field is also necessary for the Alternate-Marking Method. To color the MPLS traffic and to distinguish between hop-by-hop measurement and edge-to-edge measurement, the TC for the FL is defined as follows:

- L(oss) bit is used for coloring the MPLS packets for loss measurement. Setting the bit means color 1, and unsetting the bit means color 0.
- D(elay) bit is used for coloring the MPLS packets for delay/jitter measurement. Setting the bit means color for delay measurement.
- T(ype) bit is used to indicate the measurement type. When the T bit is set to 1, that means edge-to-edge performance measurement. When the T bit is set to 0, that means hop-by-hop performance measurement.

Considering the FL is not used as a forwarding label, the repurposing of the TC for the FL is feasible and viable.

3.1. Examples for Applying Flow-ID Label in a Label Stack

Three examples of different layouts of the FL (4 octets) are illustrated as follows. Note that more examples may exist.

3.1.1. Layout of the Flow-ID Label when Applied to MPLS Transport

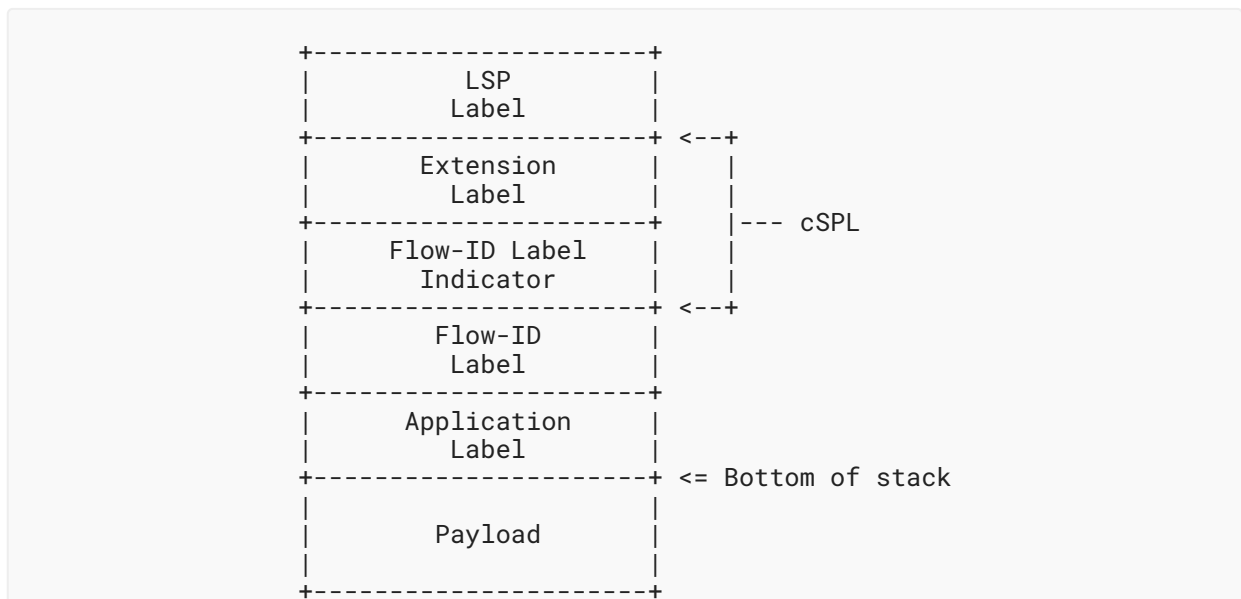


Figure 2: Applying Flow-ID to MPLS Transport

With penultimate hop popping (PHP [Section 3.16](#) of [\[RFC3031\]](#)), the top label is "popped at the penultimate LSR of the LSP, rather than at the LSP Egress". The final bullet of [Section 4](#) of the present document requires that "[t]he processing node **MUST** pop the XL, FLI, and FL from the MPLS label stack when it needs to pop the preceding forwarding label", which implies that the penultimate Label Switching Router (LSR) needs to follow the requirement of [Section 4](#) in order to support this specification. If this is done, the egress LSR is excluded from the performance measurement. Therefore, when this specification is in use, PHP should be disabled, unless the penultimate LSR is known to have the necessary support and unless it's acceptable to exclude the egress LSR.

Also note that in other examples of applying Flow-ID to MPLS transport, one LSP label can be substituted by multiple SID labels in the case of using SR Policy, and the combination of cSPL and FL can be placed between SID labels, as specified in [Section 6](#).

3.1.2. Layout of the Flow-ID Label when Applied to MPLS Service

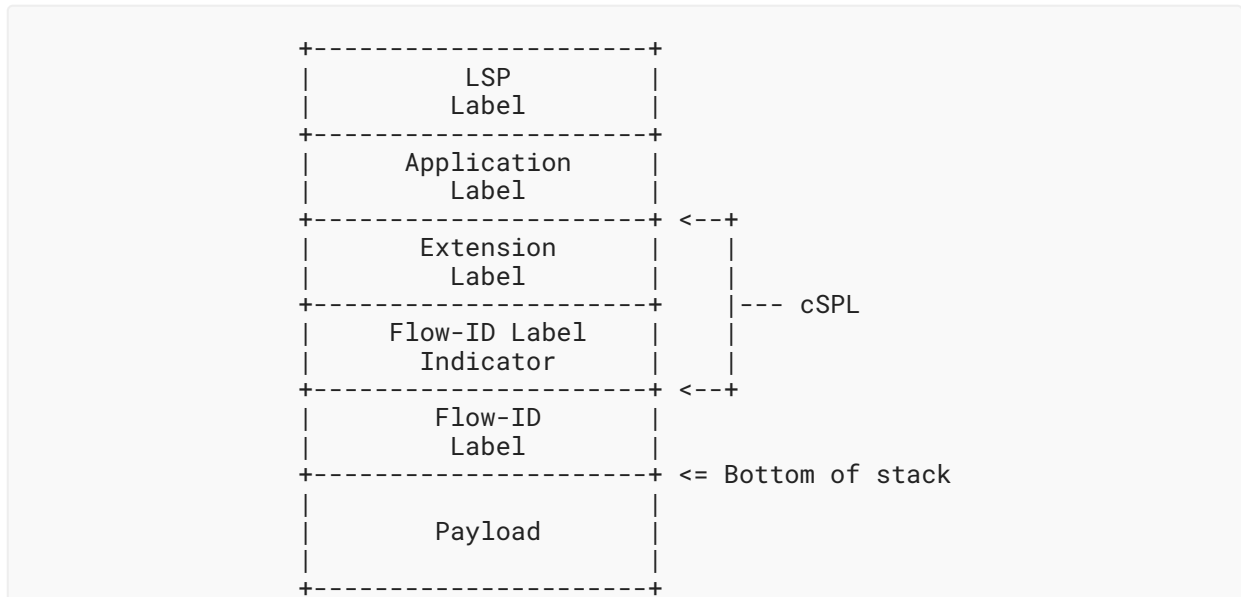


Figure 3: Applying Flow-ID to MPLS Service

Note that in this case, the application label can be an MPLS PW label, MPLS Ethernet VPN label, or MPLS IP VPN label, and it is also called a VC label as defined in [RFC4026].

3.1.3. Layout of the Flow-ID Label when Applied to both MPLS Transport and MPLS Service

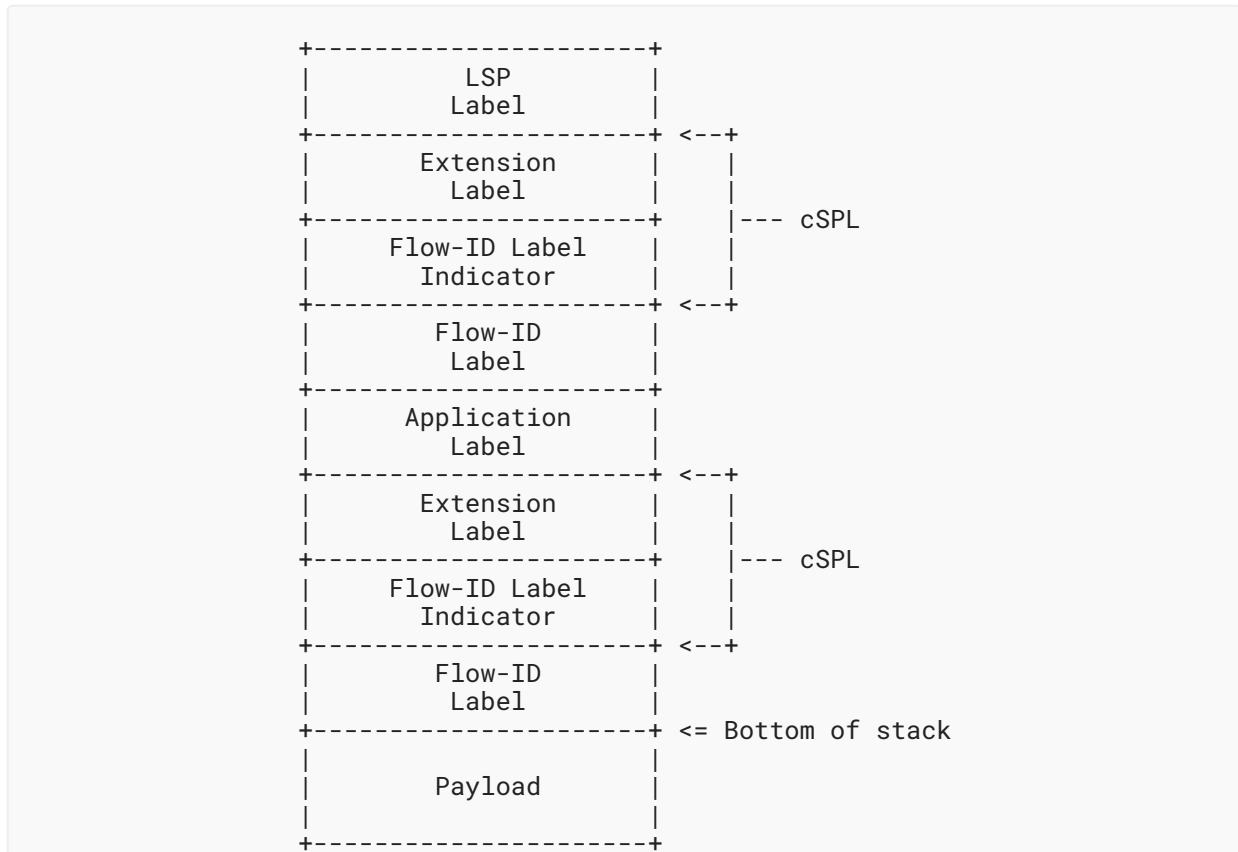


Figure 4: Applying Flow-ID to both MPLS Transport and MPLS Service

Note that for this example, the two FL values appearing in a label stack must be different. In other words, the FL applied to the MPLS transport and the FL applied to the MPLS service must be different. Also, note that the two FL values are independent of each other. For example, two packets can belong to the same VPN flow but different LSP flows, or two packets can belong to different VPN flows but the same LSP flow.

4. Procedures of Encapsulation, Look-Up, and Decapsulation

The procedures for FL encapsulation, look-up, and decapsulation are summarized as follows:

- The MPLS ingress node [RFC3031] inserts the XL, FLI, and FL into the MPLS label stack. At the same time, the ingress node sets the FL value, the two color-marking bits, and the T bit, as defined in Section 3.
- If edge-to-edge measurement is applied, i.e., the T bit is set to 1, then only the MPLS ingress/egress node [RFC3031] is the processing node; otherwise, all the MPLS nodes along the LSP are the processing nodes. The processing node looks up the FL with the help of the XL and FLI, and exports the collected data (such as the Flow-ID, block counters, and timestamps) to an external NMS/controller, referring to the Alternate-Marking Method. Section 6 of [ALTMARK] describes protocols for collected data export; the details on how to export the

collected data are outside the scope of this document. Note that while looking up the FL, the transit node needs to inspect beyond the label at the top of the label stack used to make forwarding decisions.

- The processing node **MUST** pop the XL, FLI, and FL from the MPLS label stack when it needs to pop the preceding forwarding label. The egress node **MUST** pop the whole MPLS label stack. This document doesn't introduce any new process to the decapsulated packet.

5. Procedures of Flow-ID Allocation

There are at least two ways of allocating Flow-ID. One way is to allocate Flow-ID by a manual trigger from the network operator, and the other way is to allocate Flow-ID by an automatic trigger from the ingress node. Details are as follows:

- In the case of a manual trigger, the network operator manually inputs the characteristics (e.g., IP five tuples and IP DSCP) of the measured flow; then the NMS/controller generates one or two Flow-IDs based on the input from the network operator and provisions the ingress node with the characteristics of the measured flow and the corresponding allocated Flow-ID(s).
- In the case of an automatic trigger, the ingress node identifies the flow entering the measured path and exports the characteristics of the identified flow to the NMS/controller by IPFIX [RFC7011]; then the NMS/controller generates one or two Flow-IDs based on the characteristics exported from the ingress node and provisions the ingress node with the characteristics of the identified flow and the corresponding allocated Flow-ID(s).

The policy preconfigured at the NMS/controller decides whether one Flow-ID or two Flow-IDs are generated. If the performance measurement on the MPLS service is enabled, then one Flow-ID applied to the MPLS service is generated. If the performance measurement on the MPLS transport is enabled, then one Flow-ID applied to the MPLS transport is generated. If both of them are enabled, then two Flow-IDs are respectively applied to the MPLS service and the MPLS transport are generated. In this case, a transit node needs to look up both of the two Flow-IDs by default. However, this behavior can be changed through configuration, such as by setting it to look up only the Flow-ID applied to the MPLS transport.

Whether using the two methods mentioned above or other methods to allocate Flow-ID, the NMS/controller **MUST** ensure that every generated Flow-ID is unique within the administrative domain and **MUST NOT** have any value in the reserved label space (0-15) [RFC3032]. Specifically, the statement of "Flow-ID is unique" means that the values of Flow-ID are distinct and non-redundant for any flow at any given time within an administrative domain, such that no two flows share the same Flow-ID. This uniqueness ensures that each flow can be individually identified, tracked, and differentiated from others for accurate performance monitoring and management.

6. FLC and FRLD Considerations

Analogous to the Entropy Label Capability (ELC) defined in [Section 5](#) of [RFC6790] and the Entropy Readable Label Depth (ERLD) defined in [Section 4](#) of [RFC8662], the Flow-ID Label Capability (FLC) and the Flow-ID Readable Label Depth (FRLD) are defined in this document. Both FLC and FRLD have similar semantics with the ELC and ERLD to a router, except that the Flow-ID is used in its flow identification function while the Entropy is used in its load-balancing function.

The ingress node **MUST** insert each FL at an appropriate depth, which ensures the node to which the FL is exposed has the FLC. The ingress node **SHOULD** insert each FL within an appropriate FRLD, which is the minimum FRLD of all the on-path nodes that need to read and use the FL in question. How the ingress node knows the FLC and FRLD of all the on-path nodes is outside the scope of this document.

When the SR paths are used for transport, the label stack grows as the number of on-path segments increases. If the number of on-path segments is high, that may become a challenge for the FL to be placed within an appropriate FRLD. To overcome this potential challenge, an implementation **MAY** allow the ingress node to place FL between SID labels. This means that multiple identical FLs at different depths **MAY** be interleaved with SID labels. When this occurs, sophisticated network planning may be needed, which is beyond the scope of this document.

7. Equal-Cost Multipath Considerations

Analogous to what's described in [Section 5](#) of [RFC8957], under conditions of equal-cost multipath, the introduction of the FL may lead to the same problem that is caused by the Synonymous Flow Label (SFL) [RFC8957]. The two solutions proposed for SFL also apply here. Specifically, adding FL to an existing flow may cause that flow to take a different path. If the operator expects to resolve this problem, they can choose to apply entropy labels [RFC6790] or add FL to all flows.

8. Security Considerations

As specified in [Section 7.1](#) of [RFC9341], "for security reasons, the Alternate-Marking Method **MUST** only be applied to controlled domains." This requirement applies when the MPLS performance measurement with Alternate-Marking Method is taken into account, which means the MPLS encapsulation and related procedures defined in this document **MUST** only be applied to controlled domains; otherwise, the potential attacks discussed in [Section 10](#) of [RFC9341] may be applied to the deployed MPLS networks.

As specified in [Section 3](#), the value of an FL **MUST** be unique within the administrative domain. In other words, the administrative domain is the scope of an FL. The method for achieving multi-domain performance measurement with the same FL is outside the scope of this document. The

FL **MUST NOT** be signaled and distributed outside the administrative domain. Improper configuration that allows the FL to be passed from one administrative domain to another would result in Flow-ID conflicts.

To prevent packets carrying FLs from leaking from one domain to another, domain boundary nodes **MUST** deploy policies (e.g., ACL) to filter out these packets. Specifically, at the sending edge, the domain boundary node **MUST** filter out the packets that carry the FLI and are sent to other domains. At the receiving edge, the domain boundary node **MUST** drop the packets that carry the FLI and are from other domains. Note that packet leakage is neither breaching privacy nor a source of DoS.

9. IANA Considerations

IANA has assigned the following value in the "Extended Special-Purpose MPLS Label Values" registry within the "Special-Purpose Multiprotocol Label Switching (MPLS) Label Values" registry group:

Value	Description	Reference
18	Flow-ID Label Indicator (FLI)	RFC 9714

Table 1: New Extended Special-Purpose MPLS Label Value for Flow-ID Label Indicator

10. References

10.1. 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/info/rfc2119>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.
- [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>>.
- [RFC9017] Andersson, L., Kompella, K., and A. Farrel, "Special-Purpose Label Terminology", RFC 9017, DOI 10.17487/RFC9017, April 2021, <<https://www.rfc-editor.org/info/rfc9017>>.

10.2. Informative References

- [ALT-MARK] Fioccola, G., Zhu, K., Graf, T., Nilo, M., and L. Zhang, "Alternate Marking Deployment Framework", Work in Progress, Internet-Draft, draft-ietf-ippm-alt-mark-deployment-02, 9 October 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-ippm-alt-mark-deployment-02>>.
- [RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, DOI 10.17487/RFC4026, March 2005, <<https://www.rfc-editor.org/info/rfc4026>>.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/info/rfc6790>>.
- [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>>.
- [RFC8372] Bryant, S., Pignataro, C., Chen, M., Li, Z., and G. Mirsky, "MPLS Flow Identification Considerations", RFC 8372, DOI 10.17487/RFC8372, May 2018, <<https://www.rfc-editor.org/info/rfc8372>>.
- [RFC8662] Kini, S., Kompella, K., Sivabalan, S., Litkowski, S., Shakir, R., and J. Tantsura, "Entropy Label for Source Packet Routing in Networking (SPRING) Tunnels", RFC 8662, DOI 10.17487/RFC8662, December 2019, <<https://www.rfc-editor.org/info/rfc8662>>.
- [RFC8957] Bryant, S., Chen, M., Swallow, G., Sivabalan, S., and G. Mirsky, "Synonymous Flow Label Framework", RFC 8957, DOI 10.17487/RFC8957, January 2021, <<https://www.rfc-editor.org/info/rfc8957>>.
- [RFC9341] Fioccola, G., Ed., Cociglio, M., Mirsky, G., Mizrahi, T., and T. Zhou, "Alternate-Marking Method", RFC 9341, DOI 10.17487/RFC9341, December 2022, <<https://www.rfc-editor.org/info/rfc9341>>.
- [RFC9613] Bocci, M., Ed., Bryant, S., and J. Drake, "Requirements for Solutions that Support MPLS Network Actions (MNAs)", RFC 9613, DOI 10.17487/RFC9613, August 2024, <<https://www.rfc-editor.org/info/rfc9613>>.

Acknowledgements

The authors acknowledge Loa Andersson, Tarek Saad, Stewart Bryant, Rakesh Gandhi, Greg Mirsky, Aihua Liu, Shuangping Zhan, Ming Ke, Wei He, Ximing Dong, Darren Dukes, Tony Li, James Guichard, Daniele Ceccarelli, Eric Vyncke, John Scudder, Gunter van de Velde, Roman Danyliw, Warren Kumari, Murray Kucherawy, Deb Cooley, Zaheduzzaman Sarker, and Deboraha Brungard for their careful review and very helpful comments.

They also acknowledge Italo Busi and Chandrasekar Ramachandran for their insightful MPLS-RT review and constructive comments.

Additionally, the authors thank Dhruv Dhody for the English grammar review.

Contributors

Minxue Wang

China Mobile

Email: wangminxue@chinamobile.com

Wen Ye

China Mobile

Email: yewen@chinamobile.com

Authors' Addresses

Weiqiang Cheng (EDITOR)

China Mobile

Beijing

China

Email: chengweiqiang@chinamobile.com

Xiao Min (EDITOR)

ZTE Corp.

Nanjing

China

Email: xiao.min2@zte.com.cn

Tianran Zhou

Huawei

Beijing

China

Email: zhoutianran@huawei.com

Jinyou Dai

FiberHome

Wuhan

China

Email: djy@fiberhome.com**Yoav Peleg**

Broadcom

United States of America

Email: yoav.peleg@broadcom.com