

Workgroup: MPLS
Internet-Draft:
draft-farrel-mpls-forwarder-poll-response-01
Published: 6 November 2022
Intended Status: Informational
Expires: 10 May 2023
Authors: A. Farrel
Old Dog Consulting
Anonymized Responses to a Poll on MPLS Forwarder Behavior

Abstract

As part of the work on MPLS Network Actions (MNA) several questions arose concerning how existing MPLS implementations handle Special Purpose Labels (SPLs). The details of MNA protocol extensions may depend on how existing implementations may react when they see those extensions.

In order to discover what deployed implementations currently do, the MPLS working group chairs polled participants to answer specific questions. This document is intended to report anonymized answers to that poll.

It is not intended that this document should progress to publication as an RFC.

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 10 May 2023.

Copyright Notice

Copyright (c) 2022 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. Questions](#)
- [3. Anonymized Responses](#)
 - [3.1. Response 1](#)
 - [3.2. Response 2](#)
 - [3.3. Response 3](#)
 - [3.4. Response 4](#)
 - [3.5. Response 5](#)
 - [3.6. Response 6](#)
 - [3.7. Response 7](#)
- [4. Security Considerations](#)
- [5. IANA Considerations](#)
- [6. References](#)
 - [6.1. Informative References](#)
 - [6.2. URL References](#)
- [Author's Address](#)

1. Introduction

MPLS Network Actions (MNAs) indicate actions for Label Switched Paths (LSPs) and/or MPLS packets and to transfer data needed for these actions [[I-D.ietf-mpls-mna-fwk](#)].

Various proposals have been made for how MNAs and the associated data may be encoded within MPLS packets, and these depend on the use of a new Special Purpose Label (SPL) [[RFC9017](#)].

The details of MNA protocol extensions may depend on how existing implementations may react when they see those extensions. In particular, how base SPLs (bSPLs) and extended SPLs (eSPLs) are processed when they are present in an MPLS label stack processed by an MPLS router. Furthermore, questions arose about the processing of the Time to Live (TTL) [[RFC3032](#)] and the Traffic Class (TC) field [[RFC5462](#)] of the Explicit Label Indicator (ELI) and Explicit Label (EL) Label Stack Entries (LSEs) [[RFC6790](#)].

In order to discover what deployed implementations currently do, the MPLS working group chairs polled participants to answer specific questions [[URL-poll](#)]. This document is intended to report anonymized answers to that poll.

This document is presented as a snap-shot of information. It is possible that implementations will be modified in future, or that the poll responses reported here were not accurate. Therefore, beyond acting as information to be input to the working group, this document is not intended to advance further.

2. Questions

The questions asked in the poll were as follows:

1. Does your implementation look at anything more than the top label in the label stack? If so, does it:
 - a) "scan ahead" examining the labels,
 - b) Simply use the Label Stack Entries as input to a hash,
 - c) Just search for the Bottom of Stack?
 - d) Something else.
2. In the case where your implementation looks at label values below Top of Stack:
 - a) Does the scan-ahead recognise SPLs.
 - b) If so, what does it do if the label value is an SPL (bSPL or eSPL) that it does not recognise?

(Note that this question applies to [[RFC3031](#)]/[[RFC3032](#)] implementations as well as [[RFC6790](#)]/[[RFC8662](#)] implementations.

3. What value does your implementation set as:
 - a) The ELI TC field
 - b) The ELI TTL
 - c) The EL TC field
 - d) The EL TTL

In each case what happens if the received bits in those fields are not as expected?

4. Penultimate Hop Pop

- a) How does your Penultimate Hop Pop implementation ([\[RFC3031\]](#)/[\[RFC3032\]](#)/[\[RFC3270\]](#)/[\[RFC3443\]](#)) process the TTL and TC (as EXP) from the popped Label Stack Entry?
- b) In particular, does it copy either field into the exposed top-of-stack Label Stack Entry (in the case where the popped label was not bottom of stack)?

5. Does your Penultimate Hop Pop implementation examine the exposed top-of-stack label to see whether it is a bSPL? If so, what does it do?

3. Anonymized Responses

Responses were collected over the period from the initial email until 26th October 2022.

Six responses were received and are reported here. One response reported two separate implementations which are shown separately, below.

3.1. Response 1

Answers are summarised as follows:

- 1. d) Only top label examined.
- 2. Scan ahead (only for hash) does not recognise SPLs.
- 3. No ELI/EL support.
- 4. Penultimate Hop Pop
 - a) Pipe mode.
 - b) Does not copy.
- 5. No further examination.

3.2. Response 2

Answers are summarised as follows:

- 1. b) except if any EL is found in which case each ELI is used.
- 2. Below top of stack
 - a) All known SPLs are parsed.

b)

Treated as a normal label.

3. These are set in compliance with Section 4.2 of [[RFC6790](#)]. Ignore EL TTL on reception, per [[RFC6790](#)]. ELI TTL/TC are expected to be the same as the transport label.

4. Penultimate Hop Pop

a) TC bits used depending on QoS policy.

TTL is decremented.

b) The TTL of a forwarded IP packet is set to $\text{MIN}(\text{MPLS_TTL}-1, \text{IP_TTL})$, where MPLS_TTL refers to the TTL in the outermost label in the popped stack.

The TTL of a forwarded MPLS packet is set to $\text{MIN}(\text{MPLS_TTL}-1, \text{INNER_MPLS_TTL})$, where MPLS_TTL refers to the TTL in the outermost label in the popped stack and INNER_MPLS_TTL refers to the TTL in the exposed label.

5. Ignore it except for potentially overwriting the TC based on egress QoS policy.

3.3. Response 3

Answers are summarised as follows:

1. I have no idea. I don't know our implementations in detail.

2. I have no idea.

3. I have no idea.

4. Penultimate Hop Pop

I have no idea.

5. I have no idea.

3.4. Response 4

Answers are summarised as follows:

1. Default b). Some cases for c).

2. Does not look at labels below top of stack.

3. Fields set according to [[RFC6790](#)] section 4.2.

4. In uniform mode the TTL and TC are copied to the exposed LSE.
5. No further examination.

3.5. Response 5

Answers are summarised as follows:

1. a) and b) in different implementations.
2. a) and b) skip the unrecognised SPL.
3. Set values
 - a) ECI TC = 0
 - b) ELI TTL = Assign a value or copied from IP header
 - c) EL TC = 0
 - d) EL TTL = Assign a value or copied from IP headerNo check on received fields.
4. In one mode, both fields are copied. In another mode, neither field is copied.
5. No check on received fields.

3.6. Response 6

Answers are summarised as follows:

1. c)
2. Below top of stack:
 - a) All known SPLs are parsed.
 - b) Treated as a normal label.
3. Set values
 - a) Copied from LSE inserted above ELI
 - b) Copied from LSE inserted above ELI
 - c) Copied from LSE inserted above ELI
 - d) 0 by default with (unused) option to override

4. Penultimate Hop Pop

- a) The TTL and TC are used in the forwarding plane
- b) Two configuration options exist:

*If "explicit null" is configured, the TC is copied to the explicit null LSE

*If "propagate TTL" is configured, the TTL is copied to the next LSE

Otherwise, no change is made to the next LSE

- 5. Check is applied only to ELI, in which case ELI and EL are popped.

3.7. Response 7

Answers are summarised as follows:

- 1. b)
- 2. Scan ahead (only for hash) does not recognise SPLs.
- 3. No ELI/EL support.
- 4. Penultimate Hop Pop
 - a) Pipe mode.
 - b) Does not copy.
- 5. No further examination.

4. Security Considerations

Development of a solution that is not disruptive to deployed implementations is important for a stable and secure network.

5. IANA Considerations

This document makes no requests for IANA action.

6. References

6.1. Informative References

[I-D.ietf-mpls-mna-fwk] Andersson, L., Bryant, S., Bocci, M., and T. Li, "MPLS Network Actions Framework", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-fwk-02, 21 October

2022, <<https://www.ietf.org/archive/id/draft-ietf-mpls-mna-fwk-02.txt>>.

- [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>>.
- [RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, DOI 10.17487/RFC3270, May 2002, <<https://www.rfc-editor.org/info/rfc3270>>.
- [RFC3443] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, DOI 10.17487/RFC3443, January 2003, <<https://www.rfc-editor.org/info/rfc3443>>.
- [RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, DOI 10.17487/RFC5462, February 2009, <<https://www.rfc-editor.org/info/rfc5462>>.
- [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>>.
- [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>>.
- [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>>.

6.2. URL References

- [URL-poll] IETF MPLS Working Group, "Poll on MPLS forwarder characteristics", September 2022, <<https://>>

[mailarchive.ietf.org/arch/msg/mps/
lqBWUZcQnYmmvwMoN7y4d16grN0/](mailto:mailarchive.ietf.org/arch/msg/mps/lqBWUZcQnYmmvwMoN7y4d16grN0/)>.

Author's Address

Adrian Farrel
Old Dog Consulting
United Kingdom

Email: adrian@olddog.co.uk