

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 3, 2022

X. Xu
Capitalonline
H. Assarpour
Broadcom
S. Ma
Google
F. Clad
Cisco Systems, Inc.
September 1, 2021

MPLS Payload Protocol Identifier
draft-xu-mpls-payload-protocol-identifier-09

Abstract

The MPLS label stack has no explicit protocol identifier field to indicate the protocol type of the MPLS payload. This document proposes a mechanism for containing a protocol identifier field within the MPLS packet, which is useful for any new encapsulation header (e.g., INT metadata) which may need to be encapsulated with an MPLS header.

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 March 3, 2022.

Copyright Notice

Copyright (c) 2020 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
1.1.	Requirements Language	3
2.	Terminology	3
3.	Protocol Type Field	3
4.	Data Plane Processing of PIL	4
4.1.	Egress LSRs	4
4.2.	Ingress LSRs	4
4.3.	Transit LSRs	5
4.4.	Penultimate Hop LSRs	5
5.	Signaling for PIL Processing Capability	5
6.	Alternative Approaches	5
7.	Acknowledgements	6
8.	IANA Considerations	6
9.	Security Considerations	6
10.	References	6
10.1.	Normative References	6
10.2.	Informative References	6
	Authors' Addresses	7

[1.](#) Introduction

The MPLS label stack has no explicit protocol identifier field to indicate the protocol type of the MPLS payload. This document proposes a mechanism for containing a protocol identifier field within the MPLS packet, which is useful for any new encapsulation header (e.g., INT metadata) which may need to be encapsulated with an MPLS header. With this explicit protocol identifier field, there is no need any more for each new encapsulation header to deal with the notorious first nibble issue associated with MPLS individually. More specifically, there is no need to intentionally avoid the first nibble of each new encapsulation header from being 0100 (IPv4) or 0110 (IPv6) and even worsely misuse the first nibble of each new encapsulation header as an MPLS payload type field (e.g., MPLS-BIER [[RFC8296](#)]). The tacit permission of misusing the first nibble of each new encapsulation header as an MPLS payload type field would exhaust the valuable nibble space quickly. Furthermore, there is no need to insert one additional label indicating the MPLS payload type when transporting any new encapsulation header over MPLS LSPs (e.g., transporting Network Service Header (NSH) [[RFC8300](#)] over MPLS LSPs)

therefore the signalling for that additional label is not needed anymore.

To some extent, this situation is much similar to that of the MPLS reserved label space (a.k.a., the special purpose label space) [RFC7274]. Due to the concern over the scarcity of the special-purpose label space, the extended special purpose label concept is introduced accordingly. Similarly, the IETF MPLS community should take precautions on the the scarcity of the first nibble of the MPLS payload before it is too late.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Terminology

This memo makes use of the terms defined in [RFC3032].

3. Protocol Type Field

The encapsulation format for Protocol Type field is depicted as below:

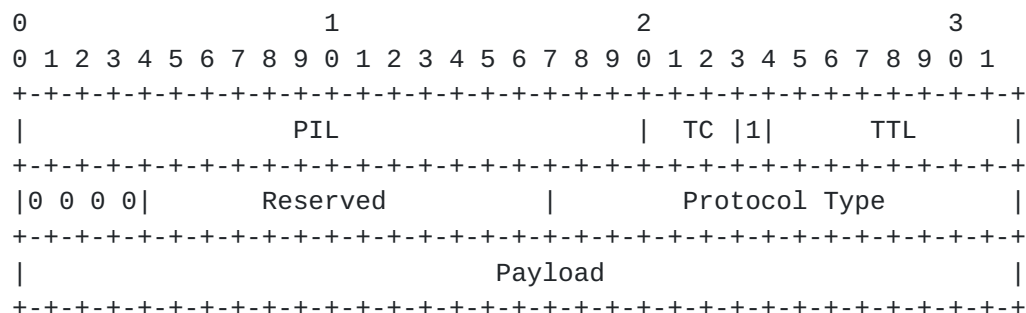


Figure 1

Protocol Identifier Label (PIL): This field contains a special purpose label with value of <TBD> or an extended special purpose label [RFC7274] with value of <TBD> which indicates that a Protocol Type field appears immediately after the bottom of the label stack.

Traffic Class (TC): The usage of this field is in accordance with the current MPLS specification [RFC3032].

S: The Bottom of Stack (BoS) field is set since the PIL MUST always appear at the bottom of the label stack.

TTL: The usage of this field is in accordance with the current MPLS specification [[RFC3032](#)].

Reserved MUST be set to 0 and ignored on reception.

Protocol Type: This field indicates the protocol type of the MPLS payload as per [[ETYPES](#)].

Payload: This field contains the MPLS payload which can be an IP packet, an Ethernet frame, or any other type of payload, e.g., Network Service Header (NSH) [[RFC8300](#)].

4. Data Plane Processing of PIL

4.1. Egress LSRs

Suppose egress LSR Y is capable of processing the Protocol Type field contained in MPLS packets. LSR Y indicates this to all ingress LSRs via signaling (see [Section 5](#)). LSR Y MUST be prepared to deal with both packets with an imposed Protocol Type field and those without; the PIL will distinguish these cases. If a particular ingress LSR chooses not to impose a Protocol Type field, LSR Y's processing of the received label stack (which might be empty) is as if LSR Y chose not to accept Protocol Type field. If an ingress LSR X chooses to impose the Protocol Type field, then LSR Y will receive an MPLS packet constructed as follows: <Top Label (TL), Application Label (AL), PIL> <Protocol Type field> <remaining MPLS payload>. Note that here the TL could be replaced with an IP-based tunnel [[RFC4023](#)] and the AL is optional. LSR Y recognizes TL as the label it distributed to its upstream LSR and pops the TL (note that the TL may be an implicit null label, in which case it doesn't appear in the label stack and LSR Y MUST process the packet starting with the AL label (if present) and/or the PIL.) LSR Y recognizes the PIL with S bit set. LSR Y then processes the Protocol Type field, which will determine how LSR Y processes the MPLS payload.

4.2. Ingress LSRs

If an egress LSR Y indicates via signaling that it can process the Protocol Type field, an ingress LSR X can choose whether or not to insert it into the MPLS packet destined for LSR Y. The ingress LSR X MUST NOT insert the Protocol Type field into that MPLS packet unless the egress LSR X has explicitly announced that it could process it. The steps that ingress LSR X performs to insert the Protocol Type field are as follows:

1. On an incoming packet, identify the application to which the packet belongs and determine whether the Protocol Type field needs to be added to the incoming packet.
2. For packets requiring the insertion of the Protocol Type field, prepend the Protocol Type field to the existing MPLS payload; then, push the PIL on to the label stack with the S bit set.
3. Push the application label (AL) label (if required) on to the label stack.
4. Push the EL and the ELI labels [[RFC6790](#)] on to the label stack (if required).
5. Determine the top label (TL) and push it on to the label stack.
6. Determine the output interface and send the packet out.

4.3. Transit LSRs

Transit LSRs MAY operate with no change in forwarding behavior. If a transit LSR recognizes the PIL and the subsequent Protocol Type field, it MAY be allowed to do some additional value-added processing, such as MPLS payload inspection, on the received MPLS packet containing the PIL and the Protocol Type field.

4.4. Penultimate Hop LSRs

No change is needed at penultimate hop LSRs.

5. Signaling for PIL Processing Capability

TBD.

6. Alternative Approaches

As illustrated in [Section 3](#) and [Section 4](#), the existence of the Protocol Type field immediately after the MPLS label stack is indicated by inserting the PIL into an MPLS packet. Alternatively, by setting the first nibble of the 4-octet entry containing the Protocol Type field to a dedicated value (e.g., 1111), the existence of the Protocol Type field could be indicated as well (see Figure 2). In this way, there is no need to insert additional label(s) (i.e., the PIL) into an MPLS packet. As for which approach should be selected in the end, it depends on a wide-scope discussion within the IETF.

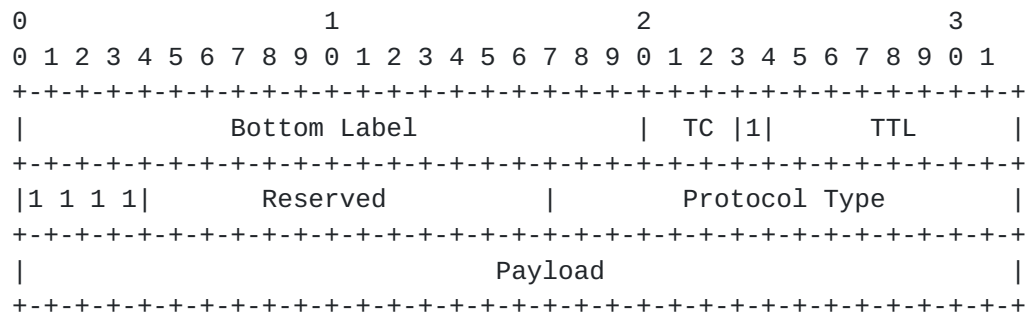


Figure 2

7. Acknowledgements

TBD.

8. IANA Considerations

A special purpose label with value of <TBD> or an extended special purpose label with value of <TBD> for the PIL needs to be assigned by the IANA

9. Security Considerations

TBD.

10. References

10.1. Normative References

- [ETYPES] The IEEE Registration Authority, "IEEE 802 Numbers", 2012.
- [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>>.

10.2. Informative References

- [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>>.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", [RFC 4023](#), DOI 10.17487/RFC4023, March 2005, <<https://www.rfc-editor.org/info/rfc4023>>.

- [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>>.
- [RFC7274] Kompella, K., Andersson, L., and A. Farrel, "Allocating and Retiring Special-Purpose MPLS Labels", [RFC 7274](#), DOI 10.17487/RFC7274, June 2014, <<https://www.rfc-editor.org/info/rfc7274>>.
- [RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", [RFC 8296](#), DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.
- [RFC8300] Quinn, P., Ed., Elzur, U., Ed., and C. Pignataro, Ed., "Network Service Header (NSH)", [RFC 8300](#), DOI 10.17487/RFC8300, January 2018, <<https://www.rfc-editor.org/info/rfc8300>>.

Authors' Addresses

Xiaohu Xu
Capitalonline.

Email: 13910161692@qq.com

Hamid Assarpour
Broadcom

Email: hamid.assarpour@broadcom.com

Shaowen Ma
Google

Email: mashaowen@gmail.com

Francois Clad
Cisco Systems, Inc.

Email: fclad@cisco.com

