

idr
Internet-Draft
Updates: 8296, 8402, 9573 (if approved)
Intended status: Standards Track
Expires: 11 December 2025

Z. Zhang, Ed.
Juniper Networks
T. Eckert, Ed.
Futurewei USA
H. Bidgoli
Nokia
Z. Zhang
ZTE
9 June 2025

BIER Payload Label
draft-zzhang-bier-payload-label-00

Abstract

The BIER Encapsulation RFC8296 specifies that the "Proto" field in the BIER header is set to 1 if an MPLS label stack follows the BIER header and the top of the stack label is a downstream-assigned label, and is set to 2 if the top of stack label is an upstream-assigned label, which is looked up in the context-specific label forwarding table that can be derived from the BIER header.

The terms upstream/downstream-assignment are inappropriate to describe the required forwarding for new MPLS label semantics such as SRGB and DCB labels as defined in RFC8402 and RFC9573 respectively. This can result in potentially inconsistent implementation choices causing potential interoperability issues.

This document rectifies this situation by changing the IANA BIER Next Protocol Identifier semantic for MPLS code points from their label semantic to the required LFIB in the forwarding plane, hence covering those new type of labels without changing behavior for existing upstream/downstream-assigned labels.

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 11 December 2025.

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	2
2. Specification	5
2.1. Update to RFC8296	5
2.2. Updates to RFC8402 and RFC9573	5
3. Security Considerations	6
4. IANA Considerations	6
5. Acknowledgments	6
6. References	6
6.1. Normative References	6
6.2. Informative References	7
Authors' Addresses	7

1. Introduction

BIER [RFC8279] is a multicast technology that achieves efficient replication without incurring per-flow/tunnel states inside a BIER domain. [RFC8296] specifies its encapsulation.

[RFC8296] specifies:

"If the BIER payload is an MPLS packet, the BIER header is followed by an MPLS label stack. ... For an example of an application where it is useful to carry an MPLS packet as the BIER payload, see [BIER_MVPN]. If the BIER encapsulation's Proto field indicates that the payload is an MPLS packet with an upstream-assigned label at the top of the stack, the upstream-assigned label is interpreted in the context of <BFIR-id, sub-domain-id>. Note that the sub-domain-id must be inferred from the BIFT-id."

The codepoints for "Proto field" in the above quote are registered in the IANA "BIER Next Protocol Identifiers" registry, with the following values for MPLS payload (among other values that are omitted here):

Value	Description	Reference
1	MPLS packet with downstream-assigned label at top of stack	[RFC8296]
2	MPLS packet with upstream-assigned label at top of stack	[RFC8296]

Table 1: Pre-existing IANA BIER Next Protocol Identifier for MPLS

MPLS downstream-assigned labels are defined in [RFC3031]. [RFC5332] introduces upstream-assigned labels as a mechanism to support multicasting of MPLS packets, such that a single MPLS packet can be received by more than one LSR. To avoid the need for all receiving LSR to coordinate the assignment of such multicasted packet, the label is assigned by the upstream (sending) LSR.

[RFC5331] specifies "Context-Specific Label Spaces" as the common MPLS forwarding plane mechanism to support upstream assigned MPLS label for multicast, as well as other use-cases that can not be supported by downstream-assigned labels, such as per-platform and per-interface label spaces defined in [RFC3031].

After [RFC8296] was published, which specifies BIER header Proto field values only for upstream and downstream labels, new MPLS label semantics where introduced: [RFC8402] introduced "SR Global Block" (SRGB), and [RFC9573] introduced "Domain-wide Control Block" (DCB). Hence, these new semantics of MPLS labels where not considered by [RFC8296] BIER Next Protocol Identifiers.

Both type of new MPLS labels can be multicast, aka: replicated and hence received by more than one receiver. In terms of [RFC5332] this could make them upstream-assigned. However, these MPLS labels are "globally" or rather domain-wide assigned, so there is no need for the use of a Context-Specific Label Space as required for upstream-assigned labels.

In the MPLS LSR forwarding plane, such Context-Specific Label Spaces are more expensive, because they require more expensive context specific lookups of MPLS labels. Therefore it is desirable that new label semantics that can use the default label space are signalled such that they will be looked up from the default label space. Both SRGB and DCB labels where designed to use the default label space.

Because the BIER Next Protocol Identifier value is meant to support (only) the necessary demultiplexing of packets in the forwarding plane on a BFER (but no further control plane semantics of the label used), it does not make sense to introduce new codepoints values for new MPLS label semantics when they are intended to share pre-existing forwarding plane semantics; specifically the default label space as used for downstream-assigned labels (codepoint 1), or a Context-Specific Label Space as used for upstream-assigned labels (codepoint 2). Additional codepoints for these cases would only complicate MPLS forwarding planes unnecessary and make introduction of new MPLS label semantics that can share these pre-existing forwarding plane demultiplexing paths unnecessary complex.

For this reason, this document changes the definitions of the two existing codepoints as follows:

- * Code point 1 indicates a label from the default label space. It is intended to be used by all label semantics that can use the default label space.
- * Code point 2 indicates a label from a context-specific label space. It is intended to be used only by those label semantics that can not use code point 1 because different sending LSR in a domain may have different bindings for the same label.

Both of these definitions are a superset of and hence backward-compatible with the existing definitions from [RFC5331], [RFC5332] and [RFC8296]. They are just not defining the use of the codepoints based on the the label semantics but on the LFIB required to support the semantic.

In addition, this document clarifies that SRGB and DCB labels are using the default label space and hence need to be signalled in the BIER header via codepoint 1 and how future semantics need to use existing or new BIER Next Protocol Identifier values.

2. Specification

2.1. Update to RFC8296

The Description for codepoint 1 and 2 in the IANA BIER Next Protocol Identifier table are changed as follows:

Value	Description	Reference
1	MPLS packet with the top of stack label to be looked up in the Default LFIB	[THISRFC]
2	MPLS packet with the top of stack label to be looked up in the Context-specific LFIB for the BFIR	[THISRFC]

Table 2: Updated IANA BIER Next Protocol Identifier Descriptions for MPLS

Downstream-assigned labels [RFC3031], "SR Global Block" (SRGB, [RFC8402]) and "Domain-wide Control Block" (DCB), [RFC9573] label use the Default LFIB for forwarding and to thus require use of value 1.

Upstream-assigned labels [RFC5331], as used by [RFC5332], [RFC8556], and [RFC9624] require the use of a context-specific LFIB and do thus require use of value 2.

New semantics for MPLS labels beyond the aforementioned (upstream, downstream, SRGB, DCB) SHOULD re-use one of these two pre-existing values unless they can not co-exist with the pre-existing MPLS label semantics in the forwarding plane. In that case, a new entry needs to be allocated in the registry.

2.2. Updates to RFC8402 and RFC9573

This document specifies that "SR Global Block" (SRGB) labels as specified in [RFC8402]) and "Domain-wide Control Block" (DCB) labels as specified in [RFC9573] have to be indicated through a value of 1 (default LFIB) when used as a top op stack label in an [RFC8296] BIER packet MPLS payload. This was underspecified in both RFCs.

3. Security Considerations

This document does not introduce new security concerns.

4. IANA Considerations

This document requests IANA to update the codepoint 1 and 2 in the "BIER Next Protocol Identifiers" registry as specified above in Table 2.

5. Acknowledgments

6. References

6.1. Normative References

- [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>>.
- [RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, DOI 10.17487/RFC5331, August 2008, <<https://www.rfc-editor.org/info/rfc5331>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [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>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC9573] Zhang, Z., Rosen, E., Lin, W., Li, Z., and IJ. Wijnands, "MVPN/EVPN Tunnel Aggregation with Common Labels", RFC 9573, DOI 10.17487/RFC9573, May 2024, <<https://www.rfc-editor.org/info/rfc9573>>.

6.2. Informative References

- [RFC5332] Eckert, T., Rosen, E., Ed., Aggarwal, R., and Y. Rekhter, "MPLS Multicast Encapsulations", RFC 5332, DOI 10.17487/RFC5332, August 2008, <<https://www.rfc-editor.org/info/rfc5332>>.
- [RFC8556] Rosen, E., Ed., Sivakumar, M., Przygienda, T., Aldrin, S., and A. Dolganow, "Multicast VPN Using Bit Index Explicit Replication (BIER)", RFC 8556, DOI 10.17487/RFC8556, April 2019, <<https://www.rfc-editor.org/info/rfc8556>>.
- [RFC9624] Zhang, Z., Przygienda, T., Sajassi, A., and J. Rabadan, "EVPN Broadcast, Unknown Unicast, or Multicast (BUM) Using Bit Index Explicit Replication (BIER)", RFC 9624, DOI 10.17487/RFC9624, August 2024, <<https://www.rfc-editor.org/info/rfc9624>>.

Authors' Addresses

Zhaohui Zhang (editor)
Juniper Networks
Email: zzhang@juniper.net

Toerless Eckert (editor)
Futurewei USA
Email: tte@cs.fau.de

Hooman Bidgoli
Nokia
Email: hooman.bidgoli@nokia.com

Zheng Zhang
ZTE
Email: zhang.zheng@zte.com.cn