

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 25 June 2026

J. Halpern
G. Mirsky
Ericsson
22 December 2025

Explicit Congestion Notification Using MPLS Network Actions
draft-halmir-mpls-ecn-01

Abstract

It has been broadly demonstrated that user experience improvements for time-critical applications have not increased at the same pace as network throughput (i.e., the increased bandwidth does not result in a corresponding increase in the user experience). Optimizing network latency rather than throughput can lead to a significantly improved user experience for time-critical applications. Low Latency, Low Loss, and Scalable Throughput (L4S) technology, introduced in RFC 9330, uses Explicit Congestion Notification (ECN) bits by marking packets suffering potential congestion in the network. L4S operates as a congestion-control mechanism, using markers within the data packets to detect and promptly respond to congestion conditions. This feedback loop enables devices (e.g., endpoints such as client devices and server devices) to adjust data flow in real-time, preventing bottlenecks and ensuring smoother transmission.

RFC 5129 describes a mechanism for supporting ECN in the Multiprotocol Label Switching (MPLS) data plane. That mechanism is based on the codepoints that take part in the Traffic Class (TC) field and, thus, prevent the use of TC field for traffic differentiation. This document defines how ECN can be supported in the MPLS data plane using the MPLS Network Actions technology.

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 25 June 2026.

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. Conventions Used in this Document	3
2.1. Acronyms	3
2.2. Requirements Language	4
3. Realization of Explicit Congestion Notification as an MPLS Network Action	4
3.1. ECN Encapsulation Using In-Stack MPLS Network Actions	4
3.2. ECN Encoding Using Post-Stack MPLS Network Actions	5
4. Theory of Operation	5
5. IANA Considerations	6
5.1. ECN as an MPLS Network Action Opcode	6
5.2. ECN as an Post-Stack Network Action Opcode	6
6. Security Considerations	7
7. Acknowledgments	7
8. References	7
8.1. Normative References	7
8.2. Informational References	8
Authors' Addresses	8

1. Introduction

Network throughput (also referred to as bandwidth) has increased exponentially over the last several decades, enabling a variety of new features for end users, including video streaming, online gaming, and mixed reality implementations, as well as more advanced applications such as remote teleoperation and autonomous services. The focus on network throughput, however, underestimates the fact that users are primarily interested in application responsiveness rather than the number of bits they can transfer per second. As a result, user experience improvements for time-critical applications

have not necessarily increased at the same pace as network throughput (i.e., the increased bandwidth does not result in a corresponding increase in the user experience). Optimizing network latency rather than throughput can lead to a significantly improved user experience for time-critical applications.

Low Latency, Low Loss, and Scalable Throughput (L4S) technology, introduced in [RFC9330], uses Explicit Congestion Notification (ECN) ([RFC3168]) bits in IP or TCP by marking packets suffering potential congestion in the network ([RFC9331]). L4S operates as a congestion-control mechanism, using ECN markers within the data packets to detect and promptly respond to congestion conditions. This feedback loop enables devices (e.g., endpoints such as client devices and server devices) to adjust data flow in real-time, preventing bottlenecks and ensuring smoother transmission.

[RFC5129] describes a mechanism for supporting ECN in the Multiprotocol Label Switching (MPLS) data plane. That mechanism is based on the codepoints that take part in the Traffic Class (TC) field [RFC3032] and, thus, prevent the use of the TC field for traffic differentiation. This document defines how ECN can be supported in the MPLS data plane using the MPLS Network Actions technology ([RFC9789]) while preserving the existing MPLS traffic class capability.

2. Conventions Used in this Document

2.1. Acronyms

ECN: Explicit Congestion Notification

ISD: In-Stack Data

L4S: Low Latency, Low Loss, and Scalable Throughput

LER: Label Edge Router

LSR: Label Switching Router

MPLS: Multiprotocol Label Switching

MNA: MPLS Network Action

NAS: Network Action Sub-stack

PSD: Post-Stack Data

TC: Traffic Class

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. Realization of Explicit Congestion Notification as an MPLS Network Action

MNA, according to [RFC9789], can support ECN by an MNA Sub-stack using either In-Stack Data (ISD) or Post-Stack Data (PSD) for ancillary data.

3.1. ECN Encapsulation Using In-Stack MPLS Network Actions

Figure 1 displays the encoding of ECN using ISD MNA.

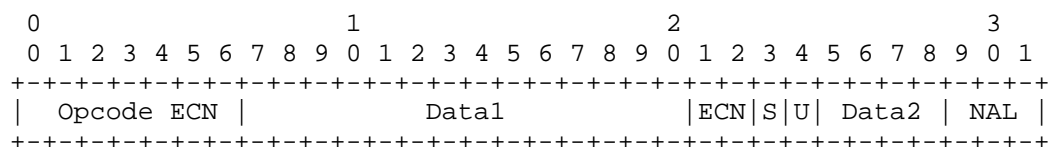


Figure 1: Format of the ECN MNA Using ISD MNA

Where fields are defined as follows:

- * Opcode ECN is seven-bit field. IANA is requested to assign (TBA1) value Section 5.1.
- * Data1 is 14-bit field. It MUST be zeroed on transmission and ignored on reception and processing of the Opcode ECN.
- * ECN is two-bit field.
- * S is the Bottom of Stack one-bit flag, as defined in [RFC3032]
- * U is a one-bit Unknown Network Action Handling field as defined in [I-D.ietf-mpls-mna-hdr].
- * Data2 is four-bit field. It MUST be zeroed on transmission and ignored on reception and processing of the Opcode ECN.
- * NAL is three-bit Network Action Length as defined in [I-D.ietf-mpls-mna-hdr].

3.2. ECN Encoding Using Post-Stack MPLS Network Actions

Figure 2 displays the encoding of ECN using PSD MNA.

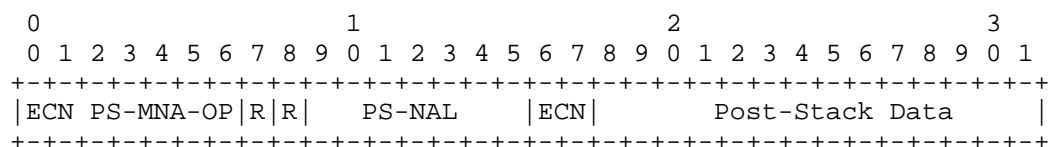


Figure 2: Format of the ECN MNA Using PSD MNA

Where the enclosed elements are defined as follows:

- * ECN PS-MNA-OP is seven-bit field. IANA is requested to assign (TBA2) value (Section 5.2).
- * R - two Reserved bits as defined in [I-D.ietf-mpls-mna-ps-hdr].
- * PS-NAL is seven-bit field as defined in [I-D.ietf-mpls-mna-ps-hdr].
- * ECN is two-bit field.
- * Post-Stack Data is 14-bit field. It MUST be zeroed on transmission and ignored on reception and processing of the ECN PS-MNA-OP.

4. Theory of Operation

If an ingress Label Edge Router (LER) that supports this specification is enabled to use ECN MNA receives a packet with the indication that the endpoints of the transport protocol are ECN-capable or the packet has experienced congestion, it MUST copy the value of the ECN field in the received IP packet into the ECN field of the ECN MNA using ISD MNA (Figure 1) or of the ECN MNA using PSD MNA (Figure 2). If the incoming packet is marked as non-ECN capable, the encapsulator SHOULD NOT use the ECN MNA. If, for some reason, the local policy requires the ECN MNA to be used, it MUST set the ECN value to Not-ECT. This corresponds to the normal mode behavior defined in [RFC6040].

If a transit Label Switching Router (LSR) that supports this specification receives an MPLS packet with ECN MNA and determines that the egress interface for the received packet experiences congestion although its buffer is not full so that the packet can be transmitted (see for the detailed explanation of the incipient congestion [RFC3168]), the LSR MUST set the ECN field to the Congestion Experienced (0b11) value.

If an egress LER that supports this specification, receives an MPLS packet with ECN MNA, it MUST follow the decapsulator's behavior defined in Section 4.2 of [RFC6040].

In some scenarios, the Readable Label Depth of the constructed path through the MPLS domain may require the placement of more than one ECN MNA Network Action Sub-stacks (NASes). If that is the case, then the Label Switching Router that exposes the NAS with ECN MNA MUST copy the value of the ECN field in that NAS into the ECN field in the next NAS that includes ECN MNA.

5. IANA Considerations

5.1. ECN as an MPLS Network Action Opcode

IANA is requested to assign an ECN codepoint (TBA1) from its Network Action Opcodes registry (creation requested in [I-D.ietf-mpls-mna-hdr]) as specified in Table 1.

Opcode	Description	Reference
TBA1	ECN as MPLS Network Action Indicator	This document

Table 1: ECN as MPLS Network Action Opcode

5.2. ECN as an Post-Stack Network Action Opcode

IANA is requested to assign an ECN codepoint (TBA1) from its Post-Stack Network Action Opcodes registry (creation requested in [I-D.ietf-mpls-mna-ps-hdr]) as specified in Table 2.

Opcode	Description	Reference
TBA2	ECN as Post-Stack Network Action Opcode	This document

Table 2: ECN as MPLS Network Action Opcode

6. Security Considerations

Security considerations discussed in [RFC9197], [RFC9326], and [RFC9789] apply to this document.

7. Acknowledgments

TBA

8. References

8.1. Normative References

- [I-D.ietf-mpls-mna-hdr]
Rajamanickam, J., Gandhi, R., Zigler, R., Song, H., and K. Kompella, "MPLS Network Action (MNA) Sub-Stack Solution", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-hdr-17, 1 December 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-hdr-17>>.
- [I-D.ietf-mpls-mna-ps-hdr]
Rajamanickam, J., Gandhi, R., Zigler, R., Li, T., and J. Dong, "Post-Stack MPLS Network Action (MNA) Solution", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-ps-hdr-04, 20 November 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-ps-hdr-04>>.
- [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>>.
- [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>>.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <<https://www.rfc-editor.org/info/rfc3168>>.
- [RFC5129] Davie, B., Briscoe, B., and J. Tay, "Explicit Congestion Marking in MPLS", RFC 5129, DOI 10.17487/RFC5129, January 2008, <<https://www.rfc-editor.org/info/rfc5129>>.

- [RFC6040] Briscoe, B., "Tunnelling of Explicit Congestion Notification", RFC 6040, DOI 10.17487/RFC6040, November 2010, <<https://www.rfc-editor.org/info/rfc6040>>.
- [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>>.
- [RFC9197] Brockners, F., Ed., Bhandari, S., Ed., and T. Mizrahi, Ed., "Data Fields for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9197, DOI 10.17487/RFC9197, May 2022, <<https://www.rfc-editor.org/info/rfc9197>>.
- [RFC9326] Song, H., Gafni, B., Brockners, F., Bhandari, S., and T. Mizrahi, "In Situ Operations, Administration, and Maintenance (IOAM) Direct Exporting", RFC 9326, DOI 10.17487/RFC9326, November 2022, <<https://www.rfc-editor.org/info/rfc9326>>.
- [RFC9789] Andersson, L., Bryant, S., Bocci, M., and T. Li, "MPLS Network Actions (MNAs) Framework", RFC 9789, DOI 10.17487/RFC9789, July 2025, <<https://www.rfc-editor.org/info/rfc9789>>.

8.2. Informational References

- [RFC9330] Briscoe, B., Ed., De Schepper, K., Bagnulo, M., and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture", RFC 9330, DOI 10.17487/RFC9330, January 2023, <<https://www.rfc-editor.org/info/rfc9330>>.
- [RFC9331] De Schepper, K. and B. Briscoe, Ed., "The Explicit Congestion Notification (ECN) Protocol for Low Latency, Low Loss, and Scalable Throughput (L4S)", RFC 9331, DOI 10.17487/RFC9331, January 2023, <<https://www.rfc-editor.org/info/rfc9331>>.

Authors' Addresses

Joel Halpern
Ericsson
Email: jmh@joelhalpern.com

Greg Mirsky
Ericsson
Email: gregimirsky@gmail.com