
Stream:	Internet Engineering Task Force (IETF)			
RFC:	9791			
Category:	Informational			
Published:	July 2025			
ISSN:	2070-1721			
Authors:	T. Saad	K. Makhijani	H. Song	G. Mirsky
	<i>Cisco Systems, Inc.</i>	<i>Independent</i>	<i>Futurewei Technologies</i>	<i>Ericsson</i>

RFC 9791

Use Cases for MPLS Network Action Indicators and Ancillary Data

Abstract

This document presents use cases that have a common feature that may be addressed by encoding network action indicators and associated ancillary data within MPLS packets. There is community interest in extending the MPLS data plane to carry such indicators and ancillary data to address these use cases.

The use cases described in this document are not an exhaustive set but rather the ones that have been actively discussed by members of the IETF MPLS, PALS, and DetNet Working Groups from the beginning of work on MPLS Network Action (MNA) until the publication of this document.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are candidates for any level of Internet Standard; see 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/rfc9791>.

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
1.1. Terminology	3
1.2. Abbreviations	3
2. Use Cases	4
2.1. No Further Fast Reroute	4
2.2. Applicability of Hybrid Measurement Methods	4
2.2.1. In Situ OAM	4
2.2.2. Alternate Marking Method	5
2.3. Network Slicing	5
2.4. NSH-Based Service Function Chaining	6
2.5. Network Programming	6
3. Coexistence with the Existing MPLS Services Using Post-Stack Headers	6
4. Coexistence of the MNA Use Cases	7
5. IANA Considerations	7
6. Security Considerations	7
7. References	8
7.1. Normative References	8
7.2. Informative References	8
Appendix A. Use Cases for Continued Discussion	10
A.1. Generic Delivery Functions	10
A.2. Delay Budgets for Time-Bound Applications	11
A.3. Stack-Based Methods for Latency Control	11
Acknowledgements	11
Contributors	12

1. Introduction

This document describes use cases that introduce functions that require special processing by forwarding hardware. The current state of the art requires allocating a new Special-Purpose Label (SPL) [RFC3032] or Extended Special-Purpose Label (eSPL). SPLs are a very limited resource, while eSPL requires an extra label stack entry per network action, which is expensive. Therefore, an MPLS Network Action (MNA) [RFC9613] approach was proposed to extend the MPLS architecture. MNA is expected to enable functions that may require carrying additional ancillary data within the MPLS packets, as well as a means to indicate that the ancillary data is present and a specific action needs to be performed on the packet.

This document lists various use cases that could benefit extensively from the MNA framework [RFC9789]. Supporting a solution of the general MNA framework provides a common foundation for future network actions that can be exercised in the MPLS data plane.

1.1. Terminology

The following terminology is used in the document:

RFC 9543 Network Slice:

Interpreted as defined in [RFC9543]. This document uses "network slice" interchangeably as a shorter version of the term "RFC 9543 Network Slice".

MPLS Ancillary Data (also referred to in this document as "ancillary data"):

Data that can be classified as:

- residing within the MPLS label stack (referred to as "in-stack data"), and
- residing after the Bottom of Stack (BoS) (referred to as "post-stack data").

1.2. Abbreviations

AMM: Alternative Marking Method

BoS: Bottom of Stack

DEX: Direct Export

eSPL: extended Special-Purpose Label

FRR: Fast Reroute

G-ACh: Generic Associated Channel

HbH: Hop by Hop

I2E: Ingress to Egress

IOAM: In situ Operations, Administration, and Maintenance

LSP: Label Switched Path

LSR: Label Switching Router

MNA: MPLS Network Action

NRP: Network Resource Partition

NSH: Network Service Header

PW: Pseudowire

SPL: Special-Purpose Label

ToS: Top of Stack

2. Use Cases

2.1. No Further Fast Reroute

MPLS Fast Reroute [\[RFC4090\]](#) [\[RFC5286\]](#) [\[RFC7490\]](#) [\[SR-TI-LFA\]](#) is a useful and widely deployed tool for minimizing packet loss in the case of a link or node failure.

Several cases exist where, once a Fast Reroute (FRR) has taken place in an MPLS network and a packet is rerouted away from the failure, a second FRR impacts the same packet on another node and may result in traffic disruption.

In such a case, the packet impacted by multiple FRR events may continue to loop between the Label Switching Routers (LSRs) that activated FRR until the packet's TTL expires. That can lead to link congestion and further packet loss. To avoid that situation, packets that FRR has redirected will be marked using MNA to preclude further FRR processing.

2.2. Applicability of Hybrid Measurement Methods

MNA can be used to carry information essential for collecting operational information and measuring various performance metrics that reflect the experience of the packet marked by MNA. Optionally, the operational state and telemetry information collected on the LSR may be transported using MNA techniques.

2.2.1. In Situ OAM

In situ Operations, Administration, and Maintenance (IOAM), defined in [\[RFC9197\]](#) and [\[RFC9326\]](#), might be used to collect operational and telemetry information while a packet traverses a particular path in a network domain.

IOAM can run in two modes: Ingress to Egress (I2E) and Hop by Hop (HbH). In I2E mode, only the encapsulating and decapsulating nodes will process IOAM data fields. In HbH mode, the encapsulating and decapsulating nodes and intermediate IOAM-capable nodes process IOAM data fields. The IOAM data fields, defined in [RFC9197], can be used to derive the operational state of the network experienced by the packet with the IOAM Header that traversed the path through the IOAM domain.

Several IOAM Option-Types have been defined:

- Pre-allocated Trace
- Incremental Trace
- Edge-to-Edge
- Proof-of-Transit
- Direct Export (DEX)

With all IOAM Option-Types except for Direct Export (DEX), the collected information is transported in the trigger IOAM packet. In the IOAM DEX Option-Type [RFC9326], the operational state and telemetry information are collected according to a specified profile and exported in a manner and format defined by a local policy. In IOAM DEX, the user data packet is only used to trigger the IOAM data to be directly exported or locally aggregated without being carried in the IOAM trigger packets.

2.2.2. Alternate Marking Method

The Alternate Marking Method (AMM), defined in [RFC9341] and [RFC9342]), is an example of a hybrid performance measurement method [RFC7799] that can be used in the MPLS network to measure packet loss and packet delay performance metrics. [RFC8957] defines the Synonymous Flow Label framework to realize AMM in the MPLS network. The MNA is an alternative mechanism that can be used to support AMM in the MPLS network.

2.3. Network Slicing

An RFC 9543 Network Slice Service [RFC9543] provides connectivity coupled with network resource commitments and is expressed in terms of one or more connectivity constructs. Section 5 of [NS-IP-MPLS] defines a Network Resource Partition (NRP) Policy as a policy construct that enables the instantiation of mechanisms to support one or more network slice services. The packets associated with an NRP may carry a marking in their network-layer header to identify this association, which is referred to as an NRP Selector. The NRP Selector maps a packet to the associated network resources and provides the corresponding forwarding treatment onto the packet.

A router that requires the forwarding of a packet that belongs to an NRP may have to decide on the forwarding action to take based on selected next hop(s) and decide on the forwarding treatment (e.g., scheduling and drop policy) to enforce based on the associated per-hop behavior.

In this case, routers that forward traffic over a physical link shared by multiple NRPs need to identify the NRP to which the packet belongs to enforce their respective forwarding actions and treatments.

MNA technologies can signal actions for MPLS packets and carry data essential for these actions. For example, MNA can carry the NRP Selector [NS-IP-MPLS] in MPLS packets.

2.4. NSH-Based Service Function Chaining

[RFC8595] describes how Service Function Chaining can be realized in an MPLS network by emulating the Network Service Header (NSH) [RFC8300] using only MPLS label stack entries.

The approach in [RFC8595] introduces some limitations, which are discussed in [SFP-VERIF]. However, the approach can benefit from the MNA framework introduced in [RFC9789].

MNA can be used to extend NSH emulation using MPLS labels [RFC8595] to support the functionality of NSH Context Headers, whether fixed or variable length. For example, MNA could support Flow ID [RFC9263] that may be used for load-balancing among Service Function Forwarders and/or the Service Functions within the same Service Function Path.

2.5. Network Programming

In Segment Routing (SR), an ingress node steers a packet through an ordered list of instructions called "segments". Each of these instructions represents a function to be called at a specific location in the network. A function is locally defined on the node where it is executed and may range from simply moving forward in the segment list to any complex user-defined behavior.

Network Programming combines SR functions to achieve a networking objective beyond mere packet routing.

Encoding a pointer to a function and its arguments within an MPLS packet transport header may be desirable. MNA can be used to encode the FUNC::ARGs to support the functional equivalent of FUNC::ARG in Segment Routing over IPv6 as described in [RFC8986].

3. Coexistence with the Existing MPLS Services Using Post-Stack Headers

Several services can be transported over MPLS networks today. These include providing Layer 3 (L3) connectivity (e.g., for unicast and multicast L3 services) and Layer 2 (L2) connectivity (e.g., for unicast PWs, multicast E-Tree, and broadcast Ethernet LAN (E-LAN) L2 services). In those cases, the user service traffic is encapsulated as the payload in MPLS packets.

For L2 service traffic, it is possible to use a Control Word (CW) [RFC4385] [RFC5085] immediately after the MPLS header to disambiguate the type of MPLS payload, prevent possible packet misordering, and allow for fragmentation. In this case, the first nibble of the data that immediately follows the MPLS BoS is set to 0b0000 to identify the presence of the PW CW.

In addition to providing connectivity to user traffic, MPLS may also transport OAM data (e.g., over MPLS Generic Associated Channels (G-AChs) [RFC5586]). In this case, the first nibble of the data that immediately follows the MPLS BoS is set to 0b0001. It indicates the presence of a control channel associated with a PW, LSP, or section.

Bit Index Explicit Replication (BIER) [RFC8296] traffic can also be encapsulated over MPLS. In this case, BIER has defined 0b0101 as the value for the first nibble of the data that immediately appears after the BoS for any BIER-encapsulated packet over MPLS.

For PWs, the G-ACh [RFC7212] uses the first four bits of the PW control word to provide the initial discrimination between data packets and packets belonging to the associated channel, as described in [RFC4385].

MPLS can be used as the data plane for Deterministic Networking (DetNet) [RFC8655]. The DetNet sub-layers, forwarding, and service are realized using the MPLS label stack, the DetNet control word [RFC8964], and the DetNet Associated Channel Header [RFC9546].

MNA-based solutions for the use cases described in this document and proposed in the future are expected to allow for coexistence and backward compatibility with all existing MPLS services.

4. Coexistence of the MNA Use Cases

Two or more of the discussed cases may coexist in the same packet. That may require the presence of multiple ancillary data (whether in-stack or post-stack ancillary data) to be present in the same MPLS packet.

For example, IOAM may provide essential functions along with network slicing to help ensure that critical network slice Service Level Objectives (SLOs) are being met by the network provider. In this case, IOAM can collect key performance measurement parameters of a network slice traffic flow as it traverses the transport network.

5. IANA Considerations

This document has no IANA actions.

6. Security Considerations

Section 7 of [RFC9789] outlines security considerations for documents that do not specify protocols. The authors have verified that these considerations are fully applicable to this document.

In-depth security analysis for each specific use case is beyond the scope of this document and will be addressed in future solution documents. It is strongly recommended that these solution documents undergo review by a security expert early in their development, ideally during the Working Group Last Call phase.

7. References

7.1. Normative References

- [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>>.

7.2. Informative References

- [GDF] Zhang, Z., Bonica, R., Kompella, K., and G. Mirsky, "Generic Delivery Functions", Work in Progress, Internet-Draft, draft-zzhang-intarea-generic-delivery-functions-03, 11 July 2022, <<https://datatracker.ietf.org/doc/html/draft-zzhang-intarea-generic-delivery-functions-03>>.
- [NS-IP-MPLS] Saad, T., Beeram, V., Dong, J., Halpern, J., and S. Peng, "Realizing Network Slices in IP/MPLS Networks", Work in Progress, Internet-Draft, draft-ietf-teas-ns-ip-mpls-05, 2 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-ns-ip-mpls-05>>.
- [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>>.
- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <<https://www.rfc-editor.org/info/rfc4090>>.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, DOI 10.17487/RFC4385, February 2006, <<https://www.rfc-editor.org/info/rfc4385>>.
- [RFC5085] Nadeau, T., Ed. and C. Pignataro, Ed., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, DOI 10.17487/RFC5085, December 2007, <<https://www.rfc-editor.org/info/rfc5085>>.
- [RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, DOI 10.17487/RFC5286, September 2008, <<https://www.rfc-editor.org/info/rfc5286>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, DOI 10.17487/RFC5586, June 2009, <<https://www.rfc-editor.org/info/rfc5586>>.
- [RFC7212] Frost, D., Bryant, S., and M. Bocci, "MPLS Generic Associated Channel (G-ACh) Advertisement Protocol", RFC 7212, DOI 10.17487/RFC7212, June 2014, <<https://www.rfc-editor.org/info/rfc7212>>.

-
- [RFC7490] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <<https://www.rfc-editor.org/info/rfc7490>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [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>>.
- [RFC8595] Farrel, A., Bryant, S., and J. Drake, "An MPLS-Based Forwarding Plane for Service Function Chaining", RFC 8595, DOI 10.17487/RFC8595, June 2019, <<https://www.rfc-editor.org/info/rfc8595>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.
- [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>>.
- [RFC8964] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "Deterministic Networking (DetNet) Data Plane: MPLS", RFC 8964, DOI 10.17487/RFC8964, January 2021, <<https://www.rfc-editor.org/info/rfc8964>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [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>>.
- [RFC9263] Wei, Y., Ed., Elzur, U., Majee, S., Pignataro, C., and D. Eastlake 3rd, "Network Service Header (NSH) Metadata Type 2 Variable-Length Context Headers", RFC 9263, DOI 10.17487/RFC9263, August 2022, <<https://www.rfc-editor.org/info/rfc9263>>.
- [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>>.
-

- [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>>.
- [RFC9342] Fioccola, G., Ed., Cociglio, M., Sapio, A., Sisto, R., and T. Zhou, "Clustered Alternate-Marking Method", RFC 9342, DOI 10.17487/RFC9342, December 2022, <<https://www.rfc-editor.org/info/rfc9342>>.
- [RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <<https://www.rfc-editor.org/info/rfc9543>>.
- [RFC9546] Mirsky, G., Chen, M., and B. Varga, "Operations, Administration, and Maintenance (OAM) for Deterministic Networking (DetNet) with the MPLS Data Plane", RFC 9546, DOI 10.17487/RFC9546, February 2024, <<https://www.rfc-editor.org/info/rfc9546>>.
- [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>>.
- [SFP-VERIF] Liu, Y. and G. Mirsky, "MPLS-based Service Function Path(SFP) Consistency Verification", Work in Progress, Internet-Draft, draft-lm-mpls-sfc-path-verification-03, 11 June 2022, <<https://datatracker.ietf.org/doc/html/draft-lm-mpls-sfc-path-verification-03>>.
- [SR-TI-LFA] Bashandy, A., Litkowski, S., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", Work in Progress, Internet-Draft, draft-ietf-rtgwg-segment-routing-ti-lfa-21, 12 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-segment-routing-ti-lfa-21>>.
- [SRTSN] Stein, Y(J)., "Segment Routed Time Sensitive Networking", Work in Progress, Internet-Draft, draft-stein-srtsn-01, 29 August 2021, <<https://datatracker.ietf.org/doc/html/draft-stein-srtsn-01>>.

Appendix A. Use Cases for Continued Discussion

Several use cases for which MNA can provide a viable solution have been discussed. The discussion of these aspirational cases is ongoing at the time of publication of the document.

A.1. Generic Delivery Functions

Generic Delivery Functions (GDFs), defined in [GDF], provide a new mechanism to support functions analogous to those supported through the IPv6 Extension Headers mechanism. For example, GDF can support fragmentation/reassembly functionality in the MPLS network by

using the Generic Fragmentation Header. MNA can support GDF by placing a GDF header in an MPLS packet within the post-stack data block [[RFC9789](#)]. Multiple GDF headers, organized as a list of headers, can also be present in the same MPLS packet.

A.2. Delay Budgets for Time-Bound Applications

The routers in a network can perform two distinct functions on incoming packets: forwarding (where the packet should be sent) and scheduling (when the packet should be sent). IEEE-802.1 Time-Sensitive Networking (TSN) and DetNet provide several mechanisms for scheduling under the assumption that routers are time-synchronized. The most effective mechanisms for delay minimization involve per-flow resource allocation.

Segment Routing (SR) is a forwarding paradigm that allows encoding forwarding instructions in the packet in a stack data structure rather than being programmed into the routers. The SR instructions are contained within a packet in the form of a First-In, First-Out stack, dictating the forwarding decisions of successive routers. Segment routing may be used to choose a path sufficiently short to be capable of providing bounded end-to-end latency but does not influence the queueing of individual packets in each router along that path.

When carried over the MPLS data plane, a solution is required to enable the delivery of such packets to their final destination within a given time budget. One approach to address this use case in SR over MPLS (SR-MPLS) is described in [[SRTSN](#)].

A.3. Stack-Based Methods for Latency Control

One efficient data structure for inserting local deadlines into the headers is a "stack", similar to that used in SR to carry forwarding instructions. The number of deadline values in the stack equals the number of routers the packet needs to traverse in the network, and each deadline value corresponds to a specific router. The Top of Stack (ToS) corresponds to the first router's deadline, while the MPLS BoS refers to the last. All local deadlines in the stack are later than or equal to the current time (upon which all routers agree), and times closer to the ToS are always earlier than or equal to times closer to the MPLS BoS.

The ingress router inserts the deadline stack into the packet headers; no other router needs to know the requirements of the time-bound flows. Hence, admitting a new flow only requires updating the ingress router's information base.

MPLS LSRs that expose the ToS label can also inspect the associated deadline carried in the packet (either in the MPLS stack as in-stack data or after BoS as post-stack data).

Acknowledgements

The authors gratefully acknowledge the input of the members of the MPLS Open Design Team. Also, the authors sincerely thank Loa Andersson, Xiao Min, Jie Dong, and Yaron Sheffer for their thoughtful suggestions and help in improving the document.

Contributors

Loa Anderssen

Bronze Dragon Consulting

Email: loa@pi.nu

Authors' Addresses

Tarek Saad

Cisco Systems, Inc.

Email: tsaad.net@gmail.com

Kiran Makhijani

Independent

Email: kiran.ietf@gmail.com

Haoyu Song

Futurewei Technologies

Email: haoyu.song@futurewei.com

Greg Mirsky

Ericsson

Email: gregimirsky@gmail.com