

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 1 December 2025

R. Gandhi, Ed.
Cisco Systems, Inc.
G. Mirsky, Ed.
Ericsson
T. Li
Juniper Networks
H. Song
Futurewei Technologies
B. Wen
Comcast
30 May 2025

Supporting In Situ Operations, Administration and Maintenance Using MPLS
Network Actions
draft-ietf-mpls-mna-ioam-03

Abstract

In situ Operations, Administration, and Maintenance (IOAM), defined in RFC 9197, is an on-path telemetry method to collect and record the operational state and telemetry information using, for example, Pre-allocated, Proof-of-Transit, Edge-To-Edge or Incremental IOAM Option, that can be used to calculate various performance metrics. RFC 9326 defined the IOAM Direct Export (IOAM-DEX) Option in which the operational state and telemetry information are collected according to the specified profile and exported in a manner and format defined by a local policy on each node along the path.

MPLS Network Actions (MNA) techniques are meant to indicate actions to be performed on any combination of Label Switched Paths, MPLS packets, and the node itself, and to transfer data needed for these actions. This document explores the MNA mechanisms to collect and transport the on-path operational state, and telemetry information IOAM data fields, including IOAM-DEX Option.

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 1 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	3
2. Conventions Used in this Document	3
2.1. Acronyms	3
2.2. Requirements Language	4
3. Applicability of IOAM and IOAM-DEX in an MPLS Network	4
4. Realization of IOAM and IOAM-DEX as MPLS Network Actions	5
4.1. Realization of IOAM and IOAM-DEX as Post-Stack MNA	5
4.1.1. MPLS Network Action Post-Stack Header for IOAM and IOAM-DEX	7
4.2. Realization of IOAM-DEX as MPLS Network Action ISD	8
4.2.1. IOAM-DEX Option-Type Encoding as MNA ISD	10
5. Considerations for IOAM and IOAM-DEX in MPLS Networks	12
5.1. Ingress-To-Egress Scope IOAM and IOAM-DEX Network Actions	12
5.2. Hop-By-Hop Scope IOAM and IOAM-DEX Network Actions	13
5.3. Node Capability	14
5.4. Nested MPLS Encapsulation	14
5.5. Readable Label Depth Consideration	15
6. Security Considerations	15
7. IANA Considerations	15
7.1. In-Stack MPLS Network Action Opcodes	15
7.2. Post-Stack Network Action Opcode	15
8. Appendix A: Examples	16
8.1. In-Stack and Pots-stack Network Action Processing Order	16

9. References	16
9.1. Normative References	16
9.2. Informational References	17
Acknowledgments	18
Contributors	18
Authors' Addresses	19

1. Introduction

In situ OAM (IOAM) [RFC9197] is an on-path telemetry method used to collect and record the operational state and telemetry information that can be used to calculate various performance metrics. Several IOAM Option types (e.g., Pre-allocated, Proof-of-Transit, Edge-To-Edge, and Incremental) use the user packet to collect the operational state and telemetry information. Such mechanisms transport the collected information to an IOAM decapsulating node (typically located at the edge of the IOAM domain within the data packet). IOAM Direct Export (IOAM-DEX) [RFC9326] is an IOAM Option type. In IOAM-DEX, the operational state and telemetry information are collected according to the specified profile and exported in a manner and format defined by a local policy on each node along the path.

MPLS Network Actions (MNA) techniques [I-D.ietf-mpls-mna-fwk] indicate actions to be performed on any combination of Label Switched Paths, MPLS packets, the node itself, and also allow for the transfer of data needed for these actions. [I-D.ietf-mpls-mna-hdr] defines mechanisms for carrying the Network Action Sub-Stack (NAS) as part of the MPLS label stack, i.e., In-Stack MNA solution. [I-D.ietf-mpls-mna-ps-hdr] defines mechanisms for carrying MNA and Ancillary Data (AD) outside the MPLS label stack, i.e., as the Post-Stack MNA solution. [I-D.ietf-mpls-mna-usecases] describes various use cases that can be realized using MNA techniques, including IOAM Pre-allocated, Proof-of-Transit, Edge-To-Edge, Incremental, and IOAM-DEX Option types.

This document describes how MNA can be used for collecting and transporting on-path operational state and telemetry information using IOAM data fields for IOAM Option types, including IOAM-DEX. Specifying the mechanism for exporting the collected information in case of the IOAM-DEX Option is outside the scope of this document.

2. Conventions Used in this Document

2.1. Acronyms

OAM: Operations, Administration, and Maintenance

HBH: Hop-By-Hop

I2E: Ingress-To-Egress

IHS: Ingress-To-Egress (I2E), Hop-By-Hop (HBH) or Select Scope

IOAM: In situ OAM

IOAM-DEX: IOAM Direct Export

IOAM-DEX-ISD-MNA: IOAM Direct Export as MPLS Network Action ISD

ISD: In-Stack Data

PSD: Post-Stack Data

PSH: Post-Stack Header

LSE: Label Stack Entry

MPLS: Multiprotocol Label Switching

MNA: MPLS Network Action

NAI: Network Action Indicator

NAS: Network Action Sub-stack

NASL: Network Action Sub-stack Length

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. Applicability of IOAM and IOAM-DEX in an MPLS Network

Pre-allocated, Incremental, Proof-of-Transit, and Edge-to-Edge IOAM Option types [RFC9197] use user packets to collect and transport the operational state and telemetry information. In some environments, for example, data center networks, this technique is useful as the available bandwidth and the use of jumbo frames can accommodate the increase of the packet payload. This document defines the Post-Stack MNA ([I-D.ietf-mpls-mna-ps-hdr]) solution supporting Pre-allocated, Incremental, Proof-of-Transit, and Edge-to-Edge IOAM Option types (Section 4.1).

However, for some use cases, e.g., mobile backhaul, in which network resources are closely controlled, collecting and transporting the telemetry information within a data packet may noticeably decrease the cost-efficiency of network operations. Although the operational state and telemetry information are essential for network automation (Section 4 of [RFC8969]), its delivery is not as critical as user packets. As such, collecting and transporting the operational state and telemetry information using the management plane is a viable option for some environments. IOAM-DEX [RFC9326] is capable of collecting the IOAM data fields defined in [RFC9197]. The processing and transport of the collected information are controlled by a local policy which is outside the scope of this specification. The performance considerations discussed in Section 5 of [RFC9326] are applicable here. In this document, the realizations of IOAM-DEX using the In-Stack MNA ([I-D.ietf-mpls-mna-hdr] and Post-Stack MNA ([I-D.ietf-mpls-mna-ps-hdr]) are defined in Section 4.2 and Section 4.1, respectively.

4. Realization of IOAM and IOAM-DEX as MPLS Network Actions

4.1. Realization of IOAM and IOAM-DEX as Post-Stack MNA

The presence of Post-Stack Header is indicated in an In-Stack Network Action Sub-stack as defined in [I-D.ietf-mpls-mna-ps-hdr] and is shown in Figure 1.

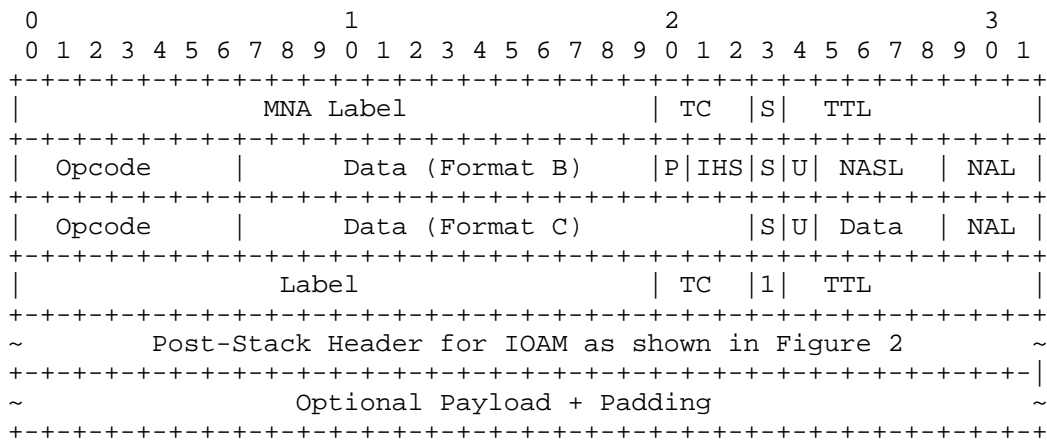


Figure 1: An Example of In-Stack Network Action Sub-Stack with Post-Stack Header Carrying IOAM and IOAM-DEX

The MNA Label (value TBA4) is defined in [I-D.ietf-mpls-mna-hdr].

The IHS scope field defined in [I-D.ietf-mpls-mna-hdr] is used to indicate that I2E or HBH or Select processing is required for the Network Action and Ancillary Data.

If both edge and intermediate nodes need to process the IOAM data fields then IHS scope MUST be set to "HBH, value 0x1". If only edge nodes need to process the IOAM data fields then IHS scope MUST be set to "I2E, value 0x0". The HBH scope allows to skip the IOAM data processing on the intermediate nodes i.e., avoids the need to parse all IOAM data fields to detect the HBH option type.

The U Flag for Unknown Action Handling is specified in [I-D.ietf-mpls-mna-hdr].

S bit is the Bottom of Stack field [RFC3032].

IHS field is set as specified in [I-D.ietf-mpls-mna-hdr].

P flag is set as specified in [I-D.ietf-mpls-mna-ps-hdr].

The Network Action Sub-Stack Length (NASL) is set as specified in [I-D.ietf-mpls-mna-hdr].

Opcode (7 bits) value is set to TBA1 (for In-Stack Network Action for IOAM and IOAM-DEX in PSH) for carrying an IOAM Option-Type and the associated data fields defined in [RFC9197] and the IOAM-DEX Option-Type and the associated data fields defined in [RFC9326] in PSH.

The In-Stack Network Action with IOAM and IOAM-DEX in PSH is optional and can be carried in Format B or Format C Label Stack Entry (LSE).

Data (10 bits, next to the Opcode field) in LSE contains the offset for PSH for this In-Stack Network Action in 4-octets units after BOS LSE to the start of the corresponding Post-Stack Network Action Opcode. Due to the Post-Stack Header type top-header, minimum value for the offset is 1 (i.e, 4-octets).

Length of Network Action (NAL) [I-D.ietf-mpls-mna-hdr] is set to 0 for the In-Stack Network Action for IOAM and IOAM-DEX option type in post-stack ancillary data.

A packet may carry more than one In-Stack Network Action in an MNA Sub-Stack for IOAM and IOAM-DEX in PSH (for example, for different IOAM Option-Types as identified in post-stack ancillary data). There may be a different In-Stack Network Action (other than for the IOAM and IOAM-DEX) in the In-Stack MNA Sub-Stack.

4.1.1. MPLS Network Action Post-Stack Header for IOAM and IOAM-DEX

An example encoding for IOAM and IOAM-DEX carried in Post-Stack Header is shown in Figure 2.

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|N N N N|Version| PS-HDR-LEN   | TYPE = MNA-POST-STACK-HDR = 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|MNA-PS-OP  |R|R|  PS-NAL    |R|R| BLOCK-NUMBER|IOAM-OPT-TYPE|
+-----+-----+-----+-----+-----+-----+-----+-----+
~   IOAM Option and Data Space [RFC9197] [RFC9326]   ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 2: An Example of Post-Stack Header Carrying IOAM and IOAM-DEX

The Post-Stack Header (PSH) type top-header is added with the following fields as defined in [I-D.ietf-mpls-mna-ps-hdr].

NNNN: The value of the first nibble as defined in [I-D.ietf-mpls-mna-ps-hdr].

PS-HDR-LEN (8 bits): PSH total length in four-octet units. This excludes the PSH type top-header.

TYPE: TYPE is defined for MNA-POST-STACK-HDR (value 1) in [I-D.ietf-mpls-mna-ps-hdr].

The PSH contains Post-Stack Network Action Opcode for IOAM and IOAM-DEX, length in number of 4-octet units, and IOAM Option-Type with IOAM data fields in the post-stack ancillary data as shown in Figure 2. The IOAM data fields MUST follow the definitions corresponding to their IOAM-Option-Types (e.g., see Section 4.4 of [RFC9197] and IOAM-DEX Option-Type in Section 3 of [RFC9326]).

An IOAM option is added in the PSH containing the following fields:

MNA-PS-OP: IANA allocated value TBA3 (Post-Stack Network Action Opcode for IOAM and IOAM-DEX) for IOAM Option-Type defined in [RFC9197], and IOAM-DEX Option-Type defined [RFC9326].

IOAM-OPT-TYPE: 7-bit field defining the IOAM Option-Type, as defined in the "IOAM Option-Type Registry" specified in [RFC9197] and [RFC9326]).

PS-NAL: 7-bit unsigned integer. Length of the IOAM data fields in

4-octet units for the MNA-PS-OP. This excludes the first 4-octet unit.

IOAM Option and Data Space: IOAM data fields as specified by the IOAM-OPT-Type field. IOAM data fields are defined corresponding to the IOAM-Option-Type (e.g., see Section 4.4 of [RFC9197] and Section 3 of [RFC9326].

BLOCK-NUMBER: The Block Number for alternate marking method can be used to aggregate the IOAM data collected in data plane, e.g., to compute measurement metrics for each block of a data flow. It is also used to correlate the IOAM data on different nodes.

4.1.1.1. Multiple IOAM and IOAM-DEX Option-Types in PSH

An example of multiple post-stack network actions with the same scope carrying different IOAM option types is shown in Figure 3.

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|N N N N|Version| PS-HDR-LEN   | TYPE = MNA-POST-STACK-HDR = 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|MNA-PS-OP |R|R|  PS-NAL     |R|R| BLOCK-NUMBER|IOAM-OPT-TYPE|
+-----+-----+-----+-----+-----+-----+-----+-----+
~   IOAM Option and Data Space [RFC9197] [RFC9326]   ~
+-----+-----+-----+-----+-----+-----+-----+-----+
.
.
.
+-----+-----+-----+-----+-----+-----+-----+-----+
|MNA-PS-OP |R|R|  PS-NAL     |R|R| BLOCK-NUMBER|IOAM-OPT-TYPE|
+-----+-----+-----+-----+-----+-----+-----+-----+
~   IOAM Option and Data Space [RFC9197] [RFC9326]   ~
+-----+-----+-----+-----+-----+-----+-----+-----+
~                               Optional Payload + Padding                               ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 3: Example of Multiple Post-Stack Network Actions with IOAM and IOAM-DEX with Same Scope

4.2. Realization of IOAM-DEX as MPLS Network Action ISD

This document uses all the elements of the IOAM-DEX Option-Type format defined in [RFC9326] to support IOAM-DEX in an MPLS network using MNA ISD [I-D.ietf-mpls-mna-hdr].

To support the direct export of the operational state and telemetry information, the IOAM-DEX-ISD-MNA blob (binary large object) is placed as part of the ISD block in an MPLS label stack according to the MNA encoding principles defined in [I-D.ietf-mpls-mna-hdr].

Using the IHS field, the IOAM-DEX-ISD-MNA can be performed in Hop-by-Hop, Ingress-to-Egress, or Select modes [I-D.ietf-mpls-mna-fwk] of collecting the operational state and telemetry information, using an MNA Opcode (Figure 4).




Figure 4: An Example of IOAM-DEX Encapsulation as an MNA Opcode

Here, the enclosed elements are defined as follows:

- * The MNA Label (value TBA4) is defined in [I-D.ietf-mpls-mna-hdr].
- * S - the Bottom of Stack field [RFC3032].
- * IHS, U, and NASL fields are set as specified in [I-D.ietf-mpls-mna-hdr].
- * P flag is set as specified in [I-D.ietf-mpls-mna-ps-hdr].
- * NASL - number of LSEs that compose the IOAM-DEX-ISD-MNA blob.
- * Opcode is MNA-IOAM-DEX opcode (value TBA2) assigned by IANA.
- * IOAM-DEX-ISD-MNA - IOAM-DEX in MPLS Network Action ISD encoding

Policies controlling the processing of the collected operational state and telemetry information, and its transport are outside the scope of this document.

4.2.1. IOAM-DEX Option-Type Encoding as MNA ISD

The IOAM-DEX in MNA header uses the LSE Format D, as defined in Section 4.4 [I-D.ietf-mpls-mna-hdr], that maps IOAM-DEX Option Type format [RFC9326]. In addition to the requirement to preserve the Bottom of Stack field, the most significant bit in LSE Format D is always set to 1 avoiding a possible mix-up of the LSE with one of the Base Special Purpose Labels. The format of IOAM-DEX-MNA blob in MNA ISD header is shown in Figure 5.

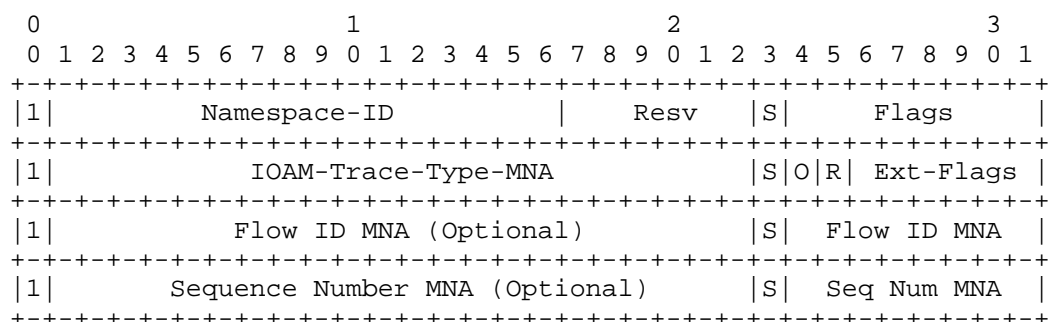


Figure 5: IOAM-DEX Option Type Format for MPLS Network Action ISD

Where fields are defined as follows:

- * Namespace-ID is a 16-bit identifier of the IOAM Namespace, as defined in [RFC9197].
- * Resv field is Reserved, it MUST be set to 0 on transmit and ignored on receipt.
- * S is a one-bit the Bottom of Stack [RFC3032].
- * Flags is an eight-bit field comprised of eight one-bit subfields. The subfields in the Flags field are allocated by IANA in IOAM-DEX Flags registry, as defined in Section 4.2 of [RFC9326].
- * IOAM-Trace-Type-MNA is a 22-bit field. The interpretation of bit positions in the IOAM-Trace-Type-MNA are as specified in [RFC9197], from Bit 0 through Bit 21. Note that bits as specified in [RFC9197] are right-shifted one bit, so for example, bit 0 in [RFC9197] (left most bit for hop_Lim and node_id in short format) is mapped to Bit 1 (second left most bit) in the MNA encoding. The registry that contains the assigned codepoints is found in the IANA's IOAM Trace-Type registry [IANA-IOAM-Trace-Type].

- * O is a one-bit flag identical to the interpretation of Bit 22 marked as "variable-length Opaque State Snapshot" as defined in [RFC9197] and assigned codepoints found in the IANA's IOAM Trace-Type registry [IANA-IOAM-Trace-Type].
- * R (Reserved) is a one-bit flag identical to the interpretation of Bit 23 marked as "reserved" as defined in [RFC9197] and assigned codepoints found in IANA's IOAM Trace-Type registry [IANA-IOAM-Trace-Type]. Similarly to [RFC9197], the R bit is reserved to allow for future extensions of the IOAM-Trace-Type-MNA bit field, i.e., it indicates the presence of the extended IOAM-Trace-Type-MNA field in the next LSE in Format D.
- * The concatenation of the 22-bit IOAM-Trace-Type-MNA, O, and R fields, explained above, is identical to IOAM-Trace-Type in the interpretation of its bits, as defined in [RFC9197] and assigned codepoints found in IANA's IOAM Trace-Type registry [IANA-IOAM-Trace-Type].
- * Ext-Flags is a six-bit field comprised of six one-bit subfields. The allocation of the subfields in the Ext-Flags field is according to Section 4.3 of [RFC9326]. The allocated flags indicate the presence of the optional Flow ID MNA and/or Sequence Number MNA fields in the IOAM-DEX-ISD-MNA header. The length of the Ext-Flags field in IOAM-DEX Option-Type in MNA is shorter by two one-bit fields compared to the length of the Extension Flags field defined in Section 3.2 of [RFC9326]. Mapping of these two bit positions are for further study. Figure 6 displays the detailed format of the Ext-Flags field.
- * Optional fields, i.e., Flow ID MNA and Sequence Number MNA, according to [RFC9326], immediately follow the Reserved field used to align optional fields at the four-octet word boundary. In the case of IOAM-DEX in MNA, such alignment can be achieved without using padding.
- * Flow ID MNA is an optional four-octet field and carries 30-bit Flow ID (after removing leading 1 and S bits). The semantics of the Flow ID MNA field is as of the Flow ID field defined in Section 3.2 of [RFC9326]. The most significant bit MUST be set to 1. Bit 23 MUST be set according to the definition of Bottom of Stack field in [RFC3032].
- * Sequence Number MNA is an optional four-octet field and carries 30-bit sequence number (after removing leading 1 and S bits). The semantics of the Sequence Number MNA field is as of the Sequence Number field defined in Section 3.2 of [RFC9326]. The most significant bit MUST be set to 1. Bit 23 MUST be set according to

the definition of Bottom of Stack field in [RFC3032]. In MPLS network environments where a label stack information is used for load-balancing flows, the 19-bit-long part of the Sequence Number MNA, starting from the Bit 1 position of the LSE, MUST remain immutable for a particular packet flow that the value of the Flow ID MNA field identifies. In MPLS networks, where other load-balancing techniques are used, all bits of the Sequence Number MNA field can be varied.

```

      0 1 2 3 4 5
    +--+--+--+--+
    |F|S|U|U|U|U|
    +--+--+--+--+

```

Figure 6: Ext-Flags Field Format

Where fields are defined as follows:

- * F - one-bit flag. When the flag is set to 1, it indicates the presence of the Flow ID field in the IOAM-DEX-ISD-MNA header.
- * S - one-bit flag. When the flag is set to 1, it indicates the presence of the Sequence Number field in the IOAM-DEX-ISD-MNA header.
- * U - unassigned one-bit flag. It MUST be zeroed on transmission and the value MUST be ignored upon receipt.

5. Considerations for IOAM and IOAM-DEX in MPLS Networks

5.1. Ingress-To-Egress Scope IOAM and IOAM-DEX Network Actions

The I2E IOAM data fields carry the IOAM Option-Type(s) that require processing on the encapsulating and decapsulating nodes only.

The IOAM Option-Type carried can be IOAM Edge-To-Edge Option-Type (value 3) defined in [RFC9197] as well as IOAM-DEX Option-Type (value 4) defined in [RFC9326]. The I2E IOAM data fields SHOULD NOT carry any IOAM Option-Type that require IOAM processing on the intermediate nodes as it will not be processed by them when IHS scope is set to "I2E, value 0x0".

The I2E IOAM and IOAM-DEX Network Action procedure is summarized as following:

- * The encapsulating node inserts an MNA Sub-Stack with the MNA Label with IHS scope set to "I2E, value 0x0", one or more In-Stack Network Actions and one or more IOAM data fields in the MPLS packet.
- * The intermediate nodes do not process the HBH IOAM data fields.
- * The decapsulating node MAY punt the IOAM data fields from the packet with the receive timestamp to the slow path for processing. The receive timestamp is required by the various I2E OAM use-cases, including streaming telemetry. Note that the packet is not necessarily punted to the control-plane.
- * The decapsulating node processes the IOAM data fields using the procedures defined in [RFC9197] and [RFC9326]. An example of IOAM processing is to export the IOAM data fields for streaming telemetry.
- * The decapsulating node MUST remove the Network Actions and IOAM data fields from the received packet. The decapsulated packet is forwarded downstream or terminated locally similar to the regular data packets.

5.2. Hop-By-Hop Scope IOAM and IOAM-DEX Network Actions

The HBH IOAM data fields carry the Option-Type(s) that require processing at the intermediate and/or encapsulating and decapsulating nodes.

The IOAM Option-Type carried can be IOAM Pre-allocated Trace Option-Type (value 0), IOAM Incremental Trace Option-Type (value 1) and IOAM Proof of Transit (POT) Option-Type (value 2), and Edge-To-Edge Option-Type (value 3) defined in [RFC9197] as well as IOAM-DEX Option-Type (value 4) defined in [RFC9326].

Editor's note: IPv6 option is not supported for HBH IOAM Incremental Trace Option-Type (value 1). Similarly, MPLS network action is also not supported for HBH IOAM Incremental Trace Option-Type (value 1) and will be removed.

The Hop-By-Hop IOAM and IOAM-DEX Network Action procedure is summarized as following:

- * The encapsulating node inserts an MNA Sub-Stack containing MNA Label, with IHS scope set to "HBH, value 0x1", one or more In-Stack Network Actions for IOAM, and one or more IOAM data fields in the MPLS packet.

- * The intermediate node enabled with HBH IOAM function processes the data packet including the IOAM data fields as defined in [RFC9197] and [RFC9326] when the node recognizes the HBH scope in the MNA Sub-Stack.
- * The intermediate node MAY punt the IOAM data fields from the packet with the receive timestamp to the slow path for processing when the node recognizes the HBH scope. The receive timestamp is required by the various HBH OAM use-cases, including streaming telemetry. Note that the packet is not necessarily punted to the control-plane.
- * The intermediate node forwards the data packet downstream.
- * The processing on the decapsulating node is same as I2E case.

Both HBH and I2E Scope IOAM may be carried in the Post-Stack MNA in an MPLS packet. In this case, the PSH with HBH IOAM data fields MUST be added after the BOS and before the PSH with I2E IOAM data fields. This way, the RLD required to process them on the intermediate nodes is minimized.

5.3. Node Capability

The decapsulating node that needs to remove the IOAM and IOAM-DEX data fields and perform the IOAM and IOAM-DEX functions may not be capable of supporting it. The encapsulating node needs to know if the decapsulating node can support the IOAM and IOAM-DEX functions. The signaling extension for this capability exchange is outside the scope of this document.

The intermediate node that is not capable of supporting the IOAM and IOAM-DEX functions defined in this document, can simply skip the IOAM and IOAM-DEX processing.

5.4. Nested MPLS Encapsulation

When a packet is received with MPLS Encapsulated Network Action for IOAM and IOAM-DEX, the nested MPLS encapsulating node that needs to add different Network Action for IOAM and IOAM-DEX, the node MUST add a new MNA Sub-Stack with the Network Action for IOAM and IOAM-DEX as part of the new MPLS encapsulation.

5.5. Readable Label Depth Consideration

The encapsulating node needs to make sure that the IOAM and IOAM-DEX data fields in MNA are added within the Readable Label Depth (RLD) of the downstream MNA capable nodes in order for them to be able to process the IOAM and IOAM-DEX.

6. Security Considerations

Security considerations discussed in [RFC9197], [RFC9326], [I-D.ietf-mpls-mna-hdr], [I-D.ietf-mpls-mna-ps-hdr] and [I-D.ietf-mpls-mna-fwk] apply to this document.

The usage of MPLS network actions defined in this document for IOAM and IOAM-DEX is intended for deployment in a single network administrative domain. As such, it assumes that the operator enabling the IOAM and IOAM-DEX operations has previously verified the integrity of the path. Still, operators need to properly secure the IOAM and IOAM-DEX in the domain to avoid malicious configuration and use, which could include injecting malicious IOAM and IOAM-DEX packets into the domain.

7. IANA Considerations

7.1. In-Stack MPLS Network Action Opcodes

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

Opcode	Description	Reference
TBA1	In-Stack Network Action for IOAM and IOAM-DEX in PSH	This document
TBA2	In-Stack Network Action for IOAM-DEX in ISD	This document

Table 1: In-Stack MPLS Network Action Opcodes

7.2. Post-Stack Network Action Opcode

IANA is requested to assign a codepoint for the Post-Stack Network Action for IOAM and IOAM-DEX Opcode from the Post-Stack Network Action Opcodes registry (to be created by in [I-D.ietf-mpls-mna-ps-hdr]).

Value	Description	Reference
TBA3	Post-Stack Network Action for IOAM and IOAM-DEX	This document

Table 2: Post-Stack Network Action Opcode

8. Appendix A: Examples

8.1. In-Stack and Pots-stack Network Action Processing Order

The In-Stack Network Action with Post-Stack Data can be added to interleave Post-Stack network actions with In-Stack network actions. The following example shows how to process the Post-Stack NA before some of the In-Stack NAs.

0										1										2										3																																																	
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																																																
MNA Label										TC										0										TTL																																																	
Opcode=8										Ancillary Data										1										IHS										0										U										NASL=3										NAL=0									
Opcode=1										Flag-Based NAIs										0										U										NAIs										NAL=0																													
Opcode=TBA1										Post-Stack Offset=1										0										U										PS-NAI										NAL=0																													
Opcode=7										Ancillary Data										1										U										AD										NAL=0																													
TBA1										Version										PS-HDR-LEN=1										TYPE = MNA-POST-STACK-HDR = 1																																																	
PS-OP=TBA3										R										R										PS-NAL=0										POST-STACK DATA																																							

Figure 7: Post-Stack and In-Stack NA Processing Order

In this example, opcode 8 is processed first, followed by the opcode 1 for Flag-Based NAIs, then the Opcode TBA1 (which has corresponding Post-Stack Network Action opcode TBA3 at an offset of 1 word, i.e., 4 bytes from the BOS), and finally opcode 7.

9. References

9.1. Normative References

[I-D.ietf-mpls-mna-fwk]

Andersson, L., Bryant, S., Bocci, M., and T. Li, "MPLS Network Actions (MNA) Framework", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-fwk-15, 27 December 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-fwk-15>>.

[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-12, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-hdr-12>>.

[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-00, 22 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-ps-hdr-00>>.

[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>>.

[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>>.

9.2. Informational References

[I-D.ietf-mpls-mna-usecases]

Saad, T., Makhijani, K., Song, H., and G. Mirsky, "Use Cases for MPLS Network Action Indicators and MPLS Ancillary Data", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-usecases-15, 23 September 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-usecases-15>>.

[IANA-IOAM-Trace-Type]

IANA, "IOAM Trace-Type", <<https://www.iana.org/assignments/ioam/ioam.xhtml#trace-type>>.

[RFC8969] Wu, Q., Ed., Boucadair, M., Ed., Lopez, D., Xie, C., and L. Geng, "A Framework for Automating Service and Network Management with YANG", RFC 8969, DOI 10.17487/RFC8969, January 2021, <<https://www.rfc-editor.org/info/rfc8969>>.

Acknowledgments

The authors would like to thank Adrian Farrel for reviewing this document and providing review comments. The authors would also like to thank Patrick Khordoc, Sagar Soni, Shwetha Bhandari, Vengada Prasad Govindan, Tarek Saad, Stewart Bryant, Xiao Min, Jaganbabu Rajamanickam, and Cheng Li for reviewing the early version of this document. The authors would also like to thank Mach Chen, Andrew Malis, Matthew Bocci, and Nick Delregno for the MPLS-RT expert review of the early version of this document.

Contributors

The following people have substantially contributed to this document:

Zafar Ali
Cisco Systems, Inc.
Email: zali@cisco.com

Loa Andersson
Huawei Technologies
Email: loa@pi.nu

Frank Brockners
Cisco Systems, Inc.
Hansaallee 249, 3rd Floor
DUESSELDORF, NORDRHEIN-WESTFALEN 40549
Germany
Email: fbrockne@cisco.com

Voitek Kozak
Comcast
Email: Voitek_Kozak@comcast.com

Authors' Addresses

Rakesh Gandhi (editor)
Cisco Systems, Inc.
Canada
Email: rgandhi@cisco.com

Greg Mirsky (editor)
Ericsson
Email: gregimirsky@gmail.com

Tony Li
Juniper Networks
Email: tony.li@tony.li

Haoyu Song
Futurewei Technologies
United States of America
Email: haoyu.song@futurewei.com

Bin Wen
Comcast
Email: Bin_Wen@cable.comcast.com