

Multiprotocol Label Switching
Internet-Draft
Intended status: Standards Track
Expires: 18 September 2026

F. Ihle
University of Tuebingen
X. Song
ZTE Corporation
M. Menth
University of Tuebingen
17 March 2026

Signaling MNA Capabilities Using LSP Ping
draft-ihlesong-mpls-mna-signaling-02

Abstract

This document defines a mechanism for discovering MPLS Network Actions (MNA) capabilities along a Label Switched Path (LSP) using the LSP Ping echo request/reply mechanism defined in RFC 8029. The In-Stack MNA capabilities include the Readable Label Depth (RLD), the maximum sizes of differently scoped Network Action Sub-stacks (MLD_NAS), and supported In-Stack network action opcodes. The Post-Stack MNA capabilities include the maximum Post-Stack MPLS Header size (MLD_PSMH), the Readable Label Depth including the Post-Stack MPLS Header (RLD_PSMH), and supported Post-Stack network action opcodes. This mechanism allows the ingress Label Edge Router (LER) to discover MNA capabilities of each transit and egress node on the path, enabling correct construction of MPLS label stacks containing MNA network actions.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://uni-tue-kn.github.io/draft-ihlesong-mpls-mna-signaling/draft-ihlesong-mpls-mna-signaling.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ihlesong-mpls-mna-signaling/>.

Discussion of this document takes place on the Multiprotocol Label Switching Working Group mailing list (<mailto:mpls@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/mpls/>. Subscribe at <https://www.ietf.org/mailman/listinfo/mpls/>.

Source for this draft and an issue tracker can be found at <https://github.com/uni-tue-kn/draft-ihlesong-mpls-mna-signaling>.

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 18 September 2026.

Copyright Notice

Copyright (c) 2026 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	4
1.1.1. Abbreviations	4
2. Definition of MNA Capabilities	6
2.1. In-Stack MNA Capabilities	6
2.1.1. The Readable Label Depth (RLD)	6
2.1.2. Maximum NAS Sizes (MLD_NAS)	7
2.1.3. Supported In-Stack Network Action Opcodes	9
2.2. Post-Stack MNA Capabilities	10
2.2.1. Post-Stack MNA Support	10
2.2.2. Maximum Post-Stack MPLS Header Size (MLD_PSMH)	10
2.2.3. Readable Label Depth Including Post-Stack MPLS Header (RLD_PSMH)	10
2.2.4. Supported Post-Stack Network Action Opcodes	11
3. LSP Ping MNA Operation Overview	11

3.1.	MNA Capabilities Query TLV	12
3.2.	MNA Capabilities Response TLV	13
3.2.1.	RLD Sub-TLV	13
3.2.2.	MLD_NAS Sub-TLV	14
3.2.3.	Supported In-Stack Opcodes Sub-TLV	15
3.2.4.	Post-Stack MNA Capabilities Sub-TLV	15
3.2.5.	Supported Post-Stack Opcodes Sub-TLV	16
4.	Processing Rules	17
4.1.	Ingress LER (Querier)	17
4.2.	Responding Node	18
4.3.	MNA-incapable Nodes	18
5.	Example	18
6.	Security Considerations	19
7.	IANA Considerations	20
7.1.	TLV Assignments	20
7.2.	New Sub-TLV Registry	20
7.3.	Return Code Assignment	21
8.	References	21
8.1.	Normative References	21
8.2.	Informative References	21
	Authors' Addresses	23

1. Introduction

The MPLS Network Actions (MNA) framework [I-D.ietf-mpls-mna-fwk] provides a general mechanism for encoding network actions and their data in the MPLS label stack. Network actions are encoded in Network Action Sub-stacks (NAS) that are placed within (ISD) or follow after (PSD) the MPLS label stack. The MNA header encoding is defined in [I-D.ietf-mpls-mna-hdr]. To correctly construct MPLS label stacks containing network actions, the ingress LER needs to know the MNA capabilities of each node along the path. For Post-Stack MNA, the ingress LER additionally needs to discover whether nodes support Post-Stack MPLS Headers and what Post-Stack network actions they can process, as required by Section 5.3 of [I-D.ietf-mpls-mna-ps-hdr]. These capabilities include:

1. In-Stack MNA capabilities:

- * The Readable Label Depth (RLD): the number of Label Stack Entries (LSEs) a node can parse without performance impact.
- * The NAS Maximum Label Depth (MLD_NAS): the maximum supported NAS size for each scope (select, HBH, I2E).
- * The supported In-Stack network action opcodes.

2. Post-Stack MNA capabilities:

- * Whether the node supports Post-Stack MNA processing as defined in [I-D.ietf-mpls-mna-ps-hdr],
- * The maximum Post-Stack MPLS Header (PSMH) size (MLD_PSMH),
- * The RLD including the PSMH (RLD_PSMH),
- * The supported Post-Stack network action opcodes.

This document defines new TLVs for the MPLS echo request/reply messages [rfc8029] to query and report MNA capabilities. The mechanism supports both "ping" mode (querying only the egress node) and "traceroute" mode (querying all nodes along the path).

1.1. Terminology

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.

1.1.1. Abbreviations

This document makes use of the terms defined in [I-D.ietf-mpls-mna-hdr] and in [I-D.ietf-mpls-mna-fwk].

Abbreviation	Name	Description	Reference
NAS	Network Action Sub-stack	A stack of related LSEs in the MPLS stack containing network actions and ancillary data.	[rfc9789]
RLD	Readable Label Depth	The number of LSEs a node can parse.	[I-D.ietf-mpls-mna-hdr]
MLD_NAS	NAS Maximum Label	The maximum number of LSEs in a	This document

	Depth	NAS that a node can process, defined per scope.	
PSMH	Post-Stack MPLS Header	The header after the BOS carrying post-stack network actions and ancillary data.	[I-D.ietf-mpls-mna-ps-hdr]
PSD	Post-Stack Data	Network actions and data encoded after the MPLS label stack.	[I-D.ietf-mpls-mna-ps-hdr]
ISD	In-Stack Data	Network actions and data encoded within the MPLS label stack.	[I-D.ietf-mpls-mna-hdr]
MLD_PSMH	Maximum PSMH Size	The maximum PSMH size a node can process, in 4-octet units.	This document
RLD_PSMH	RLD including PSMH	The total parseable depth including label stack and PSMH, in 4-octet units.	This document

Table 1: Abbreviations.

2. Definition of MNA Capabilities

This section defines the parameters that an LSR uses to signal its MNA capabilities to the ingress LER.

2.1. In-Stack MNA Capabilities

2.1.1. The Readable Label Depth (RLD)

The Readable Label Depth (RLD) is the number of LSEs an LSR can parse without performance impact [I-D.ietf-mpls-mna-fwk]. An LSR is required to search the MPLS stack for NAS that have to be processed by the LSR. To that end, the network actions must be within the RLD of the node. For HBH-scoped network actions, the ingress LER that pushes the network actions MUST ensure that the actions are readable at each LSR on the path, i.e., that they are placed within the RLD of each node.

2.1.1.1. Example

An example for the RLD parameter is given in Figure 1. With an RLD of 5, an LSR is capable of reading labels A, B, C, D, and E but not F. An RLD of 8 is required in this example to read the entire MPLS stack.

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								
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								
	MPLS-Label = A											TC			0		TTL																						
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								
	MPLS-Label = B											TC			0		TTL																						
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								
	MPLS-Label = C											TC			0		TTL																						
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								
	MPLS-Label = D											TC			0		TTL																						
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								
	MPLS-Label = E											TC			0		TTL																						
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								
	MPLS-Label = F											TC			0		TTL																						
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								
	MPLS-Label = G											TC			0		TTL																						
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								
	MPLS-Label = H											TC			1		TTL																						
+	-	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+								

Figure 1: Example MPLS stack of 8 MPLS LSEs illustrating the concept of RLD.

2.1.2. Maximum NAS Sizes (MLD_NAS)

This section gives a motivation for signaling maximum NAS sizes and then introduces the NAS Maximum Label Depth (MLD_NAS).

2.1.2.1. Motivation

A NAS in the MNA header encoding is at least 2 LSEs and at most 17 LSEs large [I-D.ietf-mpls-mna-hdr]. At an LSR, one or more NAS, e.g., a select-scoped and a hop-by-hop-scoped NAS, are possible. With two maximum-sized NAS, an LSR is required to reserve 34 LSEs in hardware to be able to process network actions. This consumes hardware resources that may be needed to encode other LSEs, e.g., forwarding labels for SR-MPLS paths, or are not available in less capable devices.

Many use cases in the MNA framework [I-D.ietf-mpls-mna-usecases] do not require a maximum-sized NAS of 17 LSEs to encode network actions and their ancillary data. Therefore, a NAS can be up to 17 LSEs but nodes can also support smaller maximum NAS. Signaling the maximum supported NAS size to the ingress LER prevents an LSR from receiving packets with a larger NAS than it supports. This way, the allocated resources for NAS can be reduced if smaller maximum NAS are supported. More resources are available for other purposes, and hardware with a low RLD can be made MNA-capable [IhMe25].

2.1.2.2. NAS Maximum Label Depth

The maximum supported number of LSEs in a NAS that an LSR can process is referred to as NAS Maximum Label Depth (MLD_NAS) in this document. For each scope in MNA, a separate parameter for the MLD_NAS exists, called MLD_NAS_Select, MLD_NAS_HBH, and MLD_NAS_I2E.

An LSR SHOULD signal the maximum-supported size of a NAS for each scope, i.e., the parameters MLD_NAS_Select, MLD_NAS_HBH, and MLD_NAS_I2E. Those parameters include the Format A, B, C, and D LSEs from [I-D.ietf-mpls-mna-hdr] in a NAS.

Based on the signaled parameters, the ingress LER MUST ensure the following when pushing the MPLS stack and NAS on a packet:

- * The ingress LER MUST NOT push a select-scoped NAS that is larger than the signaled MLD_NAS_Select value of the node that will process the select-scoped NAS.

- * The ingress LER MUST NOT push an HBH-scoped NAS that is larger than the minimum of all signaled MLD_NAS_HBH values of all nodes on the path.
- * The ingress LER MUST NOT push an I2E-scoped NAS that is larger than the signaled MLD_NAS_I2E value of the egress node.

2.1.2.3. Example

Figure 2 illustrates the different MLD_NAS sizes in an MPLS stack that are signaled by the LSR. In this example, a select-scoped NAS has a maximum size of 4 LSEs, a hop-by-hop-scoped NAS of 7 LSEs, and an I2E-scoped NAS of 4 LSEs.

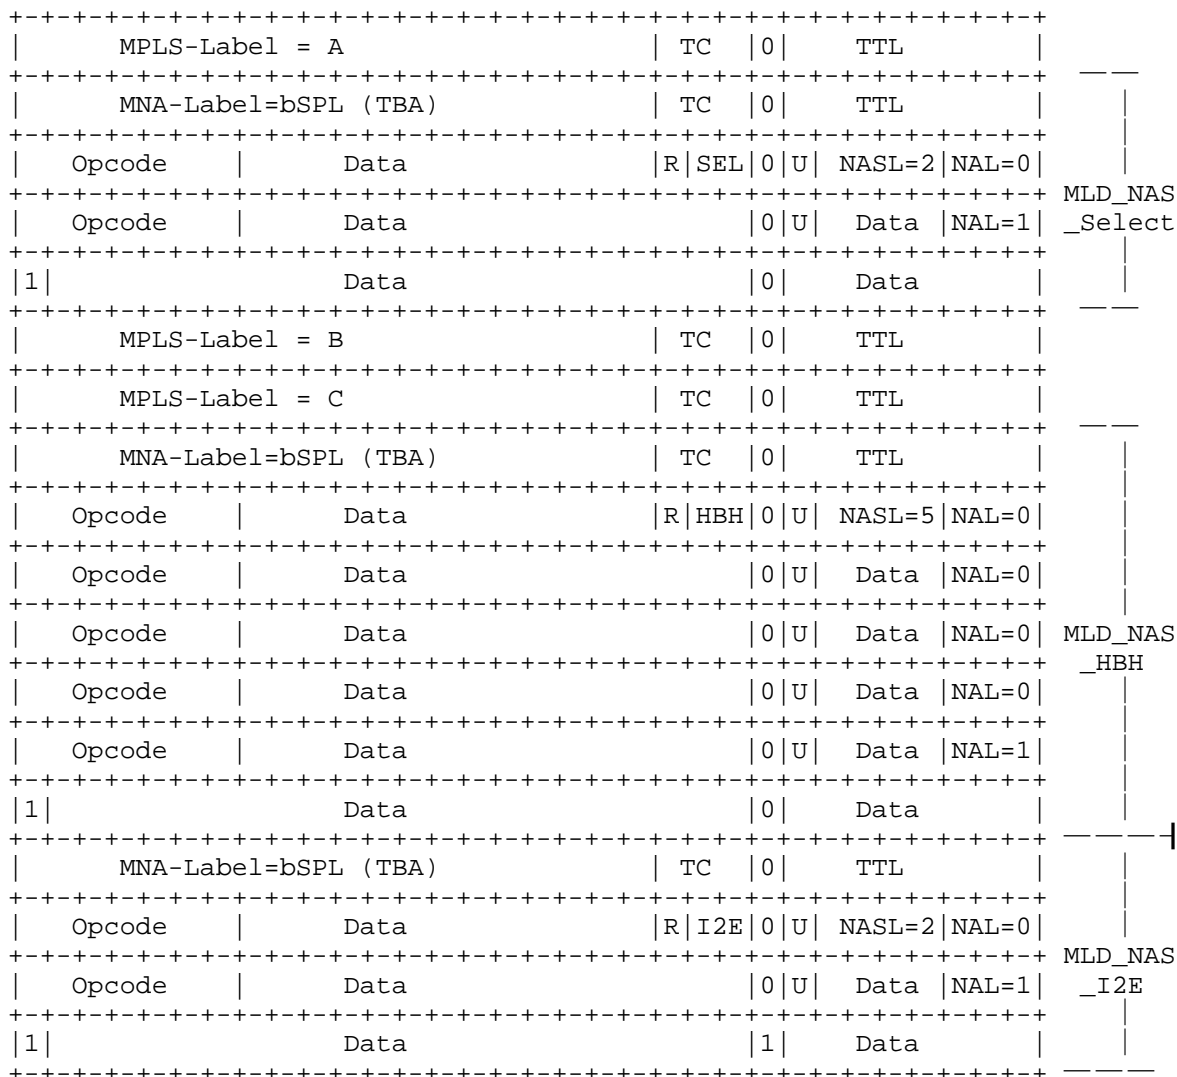


Figure 2: Example MPLS stack illustrating the different NAS sizes.

2.1.3. Supported In-Stack Network Action Opcodes

An LSR MUST signal the In-Stack network action opcodes it supports. If a network action opcode is not signaled, it is assumed that this opcode is not supported by the node.

2.2. Post-Stack MNA Capabilities

The Post-Stack MNA solution [I-D.ietf-mpls-mna-ps-hdr] allows network actions and their ancillary data to be encoded after the bottom of the MPLS label stack in a Post-Stack MPLS Header (PSMH). Section 5.3 of [I-D.ietf-mpls-mna-ps-hdr] requires that each participating node signals its Post-Stack capabilities to the encapsulating node. This section defines the parameters for that purpose.

2.2.1. Post-Stack MNA Support

A node MAY support Post-Stack MNA processing. The encapsulating node MUST NOT add a Post-Stack MPLS Header to a packet if the decapsulating node does not support Post-Stack MNA processing [I-D.ietf-mpls-mna-ps-hdr]. Therefore, the ingress LER needs to discover whether each node on the path supports Post-Stack MNA.

2.2.2. Maximum Post-Stack MPLS Header Size (MLD_PSMH)

The PSMH-LEN field in the Post-Stack MPLS Header indicates the total length of the Post-Stack MPLS Header in 4-octet units, excluding the 4-byte PSMH type header [I-D.ietf-mpls-mna-ps-hdr]. Hardware implementations may have limits on the maximum PSMH size they can process. The maximum supported PSMH length is referred to as MLD_PSMH in this document, analogous to the scope-specific values of MLD_NAS for ISD. It is expressed in 4-octet units, consistent with the PSMH-LEN field encoding. An LSR SHOULD signal its MLD_PSMH to the ingress LER. Based on the signaled parameters, the ingress LER MUST ensure the following:

- * The ingress LER MUST NOT add a PSMH with a PSMH-LEN exceeding the MLD_PSMH of any node that will process that PSMH.

2.2.3. Readable Label Depth Including Post-Stack MPLS Header (RLD_PSMH)

Section 5.3 of [I-D.ietf-mpls-mna-ps-hdr] defines the "Readable Label Depth including Post-Stack MPLS Header" as the total depth a node can parse, including both the MPLS label stack and the PSMH. This parameter is referred to as RLD_PSMH in this document and is expressed in 4-octet units. When the RLD_PSMH is signaled, the ingress LER MUST ensure that the combined size of the MPLS label stack and any PSMH intended for a node does not exceed that node's RLD_PSMH.

2.2.4. Supported Post-Stack Network Action Opcodes

The Post-Stack network action opcode space (MNA-PS-OP) is 7 bits, supporting 128 opcodes [I-D.ietf-mpls-mna-ps-hdr]. A node MUST signal the Post-Stack network action opcodes it supports. The Post-Stack opcode space is separate from the In-Stack opcode space; a node may support an opcode in-stack, post-stack, or both.

3. LSP Ping MNA Operation Overview

The MNA capability discovery mechanism operates as follows:

1. The ingress LER sends MPLS echo request messages containing the MNA Capabilities Query TLV. In traceroute mode, echo requests are sent with incrementing TTL values to reach each node on the path. In ping mode, a single echo request is sent to the egress LER.
2. Each node that receives the echo request and supports MNA capability discovery responds with an MPLS echo reply containing the MNA Capabilities Response TLV. The response includes sub-TLVs corresponding to the queried capabilities.
3. The ingress LER aggregates the received responses to determine path-wide MNA constraints. Specifically:
 - * The path-wide RLD is the minimum RLD reported by any node on the path.
 - * The path-wide MLD_NAS_HBH is the minimum MLD_NAS_HBH reported by any node.
 - * The MLD_NAS_Select for a specific node is the value reported by that node.
 - * The MLD_NAS_I2E is the value reported by the egress node.
 - * The path-wide supported opcodes for HBH-scoped NAS are the intersection of opcodes supported by all nodes.

The ingress LER SHOULD perform MNA capability discovery before pushing MNA-enabled label stacks onto a path. The ingress LER SHOULD re-query capabilities when the path changes, e.g., due to IGP reconvergence or Fast Reroute activation.

3.1. MNA Capabilities Query TLV

The MNA Capabilities Query TLV is carried in the MPLS Echo Request message. Its format is as follows:




Figure 3: MNA Capabilities Query TLV.

The fields are defined as follows:

- * Type: Indicates the MNA Capabilities Query TLV. The value is TBA1.
- * Length: The length of the Value field in octets. For this TLV, Length is 4 octets.
- * Query Flags: An 8-bit field indicating which capabilities are being queried:

Bit	Name	Description
0	QUERY_RLD	Query the Readable Label Depth
1	QUERY_MLD_NAS	Query NAS Maximum Label Depth (scopes)
2	QUERY_ISD_OPCODES	Query supported network action opcodes for ISD
3	QUERY_PS_MNA	Query Post-Stack MNA capabilities (support, MLD_PSMH, RLD_PSMH, PS opcodes)
4-7	Reserved	MUST be set to zero on transmit, ignored on receipt

Table 2: Query Flags.

- * Reserved: MUST be set to zero on transmit and MUST be ignored on receipt.

A node that receives an MNA Capabilities Query TLV with all Query Flags set to zero SHOULD respond with all available MNA capabilities.

3.2. MNA Capabilities Response TLV

The MNA Capabilities Response TLV is carried in the MPLS Echo Reply message. Its format is as follows:

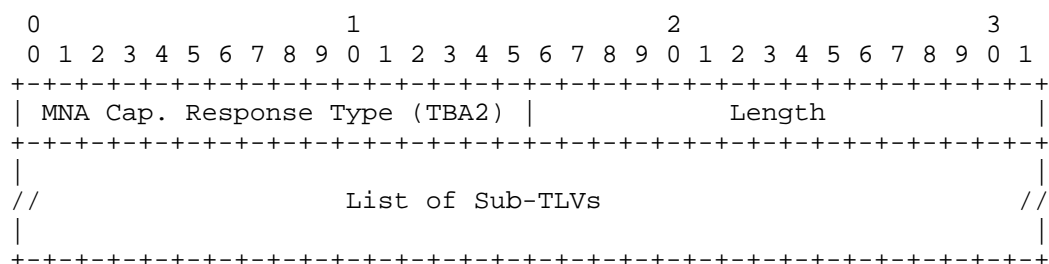


Figure 4: MNA Capabilities Response TLV.

The fields are defined as follows:

- * Type: Indicates the MNA Capabilities Response TLV. The value is TBA2.
- * Length: The length of the Value field in octets.

The Value field consists of one or more sub-TLVs as defined in the following sections. The responding node MUST include sub-TLVs corresponding to the flags set in the Query TLV. If no flags were set in the query, the responding node SHOULD include all sub-TLVs for which it has information.

3.2.1. RLD Sub-TLV

The RLD Sub-TLV reports the Readable Label Depth of the responding node.




Figure 5: RLD Sub-TLV.

- * Sub-type: 1 (RLD).
- * Length: 4 octets.
- * RLD Value: An 8-bit unsigned integer indicating the number of LSEs the node can parse without performance impact. A value of 0 indicates that the node did not provide an RLD value.
- * Reserved: MUST be set to zero on transmit and MUST be ignored on receipt.

3.2.2. MLD_NAS Sub-TLV

The MLD_NAS Sub-TLV reports the maximum supported NAS sizes for each scope. All three scope values are encoded in a single sub-TLV.




Figure 6: MLD_NAS Sub-TLV.

- * Sub-type: 2 (MLD_NAS).
- * Length: 4 octets.
- * MLD_NAS_Select: An 8-bit unsigned integer indicating the maximum number of LSEs in a select-scoped NAS that the node can process. Valid range: 2-17. A value of 0 indicates that select-scoped NAS are not supported. Values of 1 and 18-255 are invalid and MUST NOT be sent; receivers MUST treat them as 0.
- * MLD_NAS_HBH: An 8-bit unsigned integer indicating the maximum number of LSEs in an HBH-scoped NAS that the node can process. Valid range: 2-17. A value of 0 indicates that HBH-scoped NAS are not supported. Values of 1 and 18-255 are invalid and MUST NOT be sent; receivers MUST treat them as 0.

- * MLD_NAS_I2E: An 8-bit unsigned integer indicating the maximum number of LSEs in an I2E-scoped NAS that the node can process. Valid range: 2-17. A value of 0 indicates that I2E-scoped NAS are not supported. Values of 1 and 18-255 are invalid and MUST NOT be sent; receivers MUST treat them as 0.
- * Reserved: MUST be set to zero on transmit and MUST be ignored on receipt.

3.2.3. Supported In-Stack Opcodes Sub-TLV

The Supported In-Stack Opcodes Sub-TLV reports the In-Stack network action opcodes supported by the responding node using a bitmap encoding. The MNA opcode space is 7 bits, supporting 128 opcodes. Each bit in the bitmap corresponds to one opcode value.

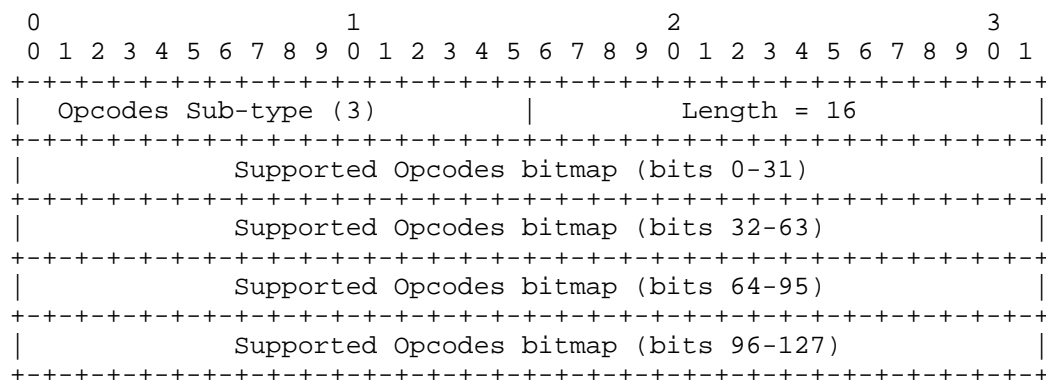


Figure 7: Supported In-Stack Opcodes Sub-TLV.

- * Sub-type: 3 (Supported ISD Opcodes).
- * Length: 16 octets.
- * Supported ISD Opcodes bitmap: A 128-bit bitmap where bit N (counting from bit 0 as the most significant bit of the first octet) corresponds to opcode value N. If bit N is set to 1, the node supports opcode N. If bit N is set to 0, the node does not support opcode N.

3.2.4. Post-Stack MNA Capabilities Sub-TLV

The Post-Stack MNA Capabilities Sub-TLV reports whether the node supports Post-Stack MNA processing, the maximum PSMH size, and the RLD including PSMH.




Figure 8: Post-Stack MNA Capabilities Sub-TLV.

- * Sub-type: 4 (Post-Stack MNA Capabilities).
- * Length: 4 octets.
- * PS Flags: An 8-bit field.
 - Bit 0: PS_SUPPORTED. If set to 1, the node supports Post-Stack MNA processing as defined in [I-D.ietf-mpls-mna-ps-hdr]. If set to 0, Post-Stack MNA is not supported and the remaining fields in this sub-TLV SHOULD be ignored.
 - Bits 1-7: Reserved. MUST be set to zero on transmit and MUST be ignored on receipt.
- * MLD_PSMH: An 8-bit unsigned integer indicating the maximum Post-Stack MPLS Header length (in 4-octet units, excluding the PSMH type header) that the node can process. A value of 0 indicates that the node did not provide this value. The valid range corresponds to the 8-bit PSMH-LEN field defined in [I-D.ietf-mpls-mna-ps-hdr].
- * RLD_PSMH: An 8-bit unsigned integer indicating the Readable Label Depth including the Post-Stack MPLS Header, in 4-octet units. A value of 0 indicates that the node did not provide this value.
- * Reserved: MUST be set to zero on transmit and MUST be ignored on receipt.

3.2.5. Supported Post-Stack Opcodes Sub-TLV

The Supported Post-Stack Opcodes Sub-TLV reports the Post-Stack network action opcodes supported by the responding node. The Post-Stack opcode space is 7 bits (128 values), identical to the In-Stack Opcodes Sub-TLV format in Section 3.2.3 but independent from it. For the Supported Post-Stack Opcodes Sub-TLV, the sub-type 5 (Supported Post-Stack Opcodes) is used. The format is identical to Figure 7.

4. Processing Rules

This section defines the processing rules for querying and responding nodes.

4.1. Ingress LER (Querier)

The ingress LER constructs MPLS echo request messages containing the MNA Capabilities Query TLV. In traceroute mode, the ingress LER sends echo requests with TTL values from 1 to the path length. This allows discovery of MNA capabilities at each hop. Traceroute mode SHOULD be used when HBH-scoped network actions are planned, as the ingress LER needs the capabilities of every node to correctly place NAS copies within each node's RLD. In ping mode, a single echo request with TTL set to 255 is sent. This is sufficient when only I2E-scoped network actions are planned, as only the egress node's capabilities are needed. After collecting responses, the ingress LER computes path-wide constraints as described in Section 3. The ingress LER MUST ensure the following when constructing MPLS stacks with MNA:

1. A select-scoped NAS pushed for a specific node MUST NOT exceed that node's MLD_NAS_Select.
2. An HBH-scoped NAS MUST NOT exceed the minimum MLD_NAS_HBH across all nodes on the path.
3. An I2E-scoped NAS MUST NOT exceed the egress node's MLD_NAS_I2E.
4. All NAS intended for a node MUST be within that node's RLD.

When Post-Stack MNA is used, the ingress LER MUST additionally ensure:

1. The ingress LER MUST NOT add a Post-Stack MPLS Header to a packet if the decapsulating node does not have the PS_SUPPORTED flag set.
2. The PSMH-LEN of any Post-Stack MPLS Header MUST NOT exceed the MLD_PSMH of any node that will process the PSMH.
3. The combined size of the MPLS label stack and any PSMH intended for a node MUST NOT exceed that node's RLD_PSMH.

4.2. Responding Node

A node that supports MNA and receives an MPLS echo request containing the MNA Capabilities Query TLV MUST respond with an MPLS echo reply containing the MNA Capabilities Response TLV. The responding node MUST include sub-TLVs corresponding to the flags set in the query. If the QUERY_RLD flag is set, the RLD Sub-TLV MUST be included. If the QUERY_MLD_NAS flag is set, the MLD_NAS Sub-TLV MUST be included. If the QUERY_ISD_OPCODES flag is set, the Supported In-Stack Opcodes Sub-TLV MUST be included. If no Query Flags are set (all zero), the responding node SHOULD include all available sub-TLVs. The reported capabilities are those of the node as a whole. If capabilities vary per interface, the node SHOULD report the capabilities applicable to the interface on which the echo request was received. If the QUERY_PS_MNA flag is set, the Post-Stack MNA Capabilities Sub-TLV (sub-type 4) MUST be included. If the node supports Post-Stack MNA and the QUERY_PS_MNA flag is set, the Supported Post-Stack Opcodes Sub-TLV (sub-type 5) MUST also be included.

4.3. MNA-incapable Nodes

A node that does not support MNA will not recognize the MNA Capabilities Query TLV. According to RFC 8029, the handling depends on the TLV type value range. The TLV type for the MNA Capabilities Query TLV SHOULD be assigned from the range that requires an error message if the TLV is not recognized. This allows the ingress LER to detect nodes that do not support MNA.

If a node does not support MNA, but recognizes the MNA Capabilities Query TLV, it MUST include Return Code TBA3 in the MPLS echo reply message.

5. Example

Consider an SR-MPLS path with three LSRs: R1, R2 (transit), and R3 (egress). The ingress LER R0 wants to push an HBH-scoped NAS and a select-scoped NAS for R2 along this path.

R0 sends MPLS echo requests in traceroute mode with all Query Flags set. The responses are:

Node	RLD	MLD_NAS_Select	MLD_NAS_HBH	MLD_NAS_I2E	PS_Supported	MLD_PSMH	RLD_PSMH
R1	20	9	9	0 (not egress)	Yes	16	36
R2	51	9	3	0 (not egress)	Yes	8	59
R3	35	9	9	9	Yes	16	51

Table 3: Example MNA Capabilities Responses.

From these responses, R0 determines:

- * Path-wide RLD: $\min(20, 51, 35) = 20$ LSEs
- * Path-wide MLD_NAS_HBH: $\min(9, 3, 9) = 3$ LSEs
- * MLD_NAS_Select for R2: 9 LSEs
- * MLD_NAS_I2E: 9 LSEs (from R3)
- * Post-Stack MNA is supported on all nodes (PS_SUPPORTED set at R1, R2, R3).
- * Path-wide MLD_PSMH for HBH-scoped PSMH: $\min(16, 8, 16) = 8$ (in 4-octet units).
- * MLD_PSMH for I2E-scoped PSMH: 16 (from R3).
- * Path-wide RLD_PSMH: $\min(36, 59, 51) = 36$ (in 4-octet units).

R0 can now construct a label stack ensuring that all NAS are within each node's RLD and do not exceed the per-scope MLD_NAS constraints. For Post-Stack MNA, R0 ensures that the PSMH does not exceed the path-wide MLD_PSMH and that the combined label stack and PSMH do not exceed any node's RLD_PSMH.

6. Security Considerations

The security considerations described in [rfc8029] apply to this document. The MNA capability discovery mechanism reveals information about node capabilities, which could potentially be exploited by an attacker to craft targeted attacks against nodes with limited MNA support. Nodes that support MNA capability discovery SHOULD support configuration options to enable or disable the MNA Capabilities

Query/Response functionality. By default, MNA capability discovery SHOULD be enabled only within an MNA-capable MPLS domain. The security considerations from [I-D.ietf-mpls-mna-hdr] and [rfc9789] also apply.

7. IANA Considerations

This section requests new TLVs and sub-TLVs.

7.1. TLV Assignments

IANA is requested to assign two new TLVs from the "TLV" registry in the "Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry group. The TLV values SHOULD be assigned from the range that requires an error message if the TLV is not recognized.

Value	TLV Name	Reference	Sub-TLV Registry
TBA1	MNA Capabilities Query	This document	No
TBA2	MNA Capabilities Response	This document	See Table 5

Table 4: New TLVs.

7.2. New Sub-TLV Registry

IANA is requested to create a new sub-TLV registry for TLV TBA2 with the following initial entries:

Sub-Type	Sub-TLV Name	Reference
0	Reserved	This document
1	RLD	This document
2	MLD_NAS	This document
3	Supported ISD Opcodes	This document
4	Post-Stack MNA Capabilities	This document
5	Supported Post-Stack Opcodes	This document

Table 5: Sub-TLV Registry for TLV TBA2.

7.3. Return Code Assignment

IANA is requested to assign a new Return Code from the "Return Code" registry in the "Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry group as follows.

Value	Meaning	Reference
TBA3	MNA not supported	This document

Table 6: New return code.

8. References

8.1. Normative References

- [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/rfc/rfc2119>>.
- [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/rfc/rfc8174>>.

8.2. Informative 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 Specification including In-Stack Network Actions and Data", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-hdr-21, 24 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-hdr-21>>.

[I-D.ietf-mpls-mna-ps-hdr]

Rajamanickam, J., Gandhi, R., Zigler, R., Dong, J., and J. Bhattacharya, "Post-Stack MPLS Network Action (MNA) Solution", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-ps-hdr-07, 24 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-ps-hdr-07>>.

[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>>.

[IhMe25]

Ihle, F. and M. Menth, "MPLS Network Actions; Technological Overview and P4-Based Implementation on a High-Speed Switching ASIC", DOI 10.1109/OJCOMS.2025.3557082, 2 April 2025, <<https://ieeexplore.ieee.org/document/10947349>>.

[rfc8029]

Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <<https://www.rfc-editor.org/rfc/rfc8029>>.

[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/rfc/rfc9789>>.

Authors' Addresses

Fabian Ihle
University of Tuebingen
Tuebingen
Germany
Email: fabian.ihle@uni-tuebingen.de

Xueyan Song
ZTE Corporation
China
Email: song.xueyan2@zte.com.cn

Michael Menth
University of Tuebingen
Tuebingen
Germany
Email: michael.menth@uni-tuebingen.de