

Workgroup: MPLS Working Group
Internet-Draft: draft-ietf-mpls-mna-hdr-04
Published: 21 October 2023
Intended Status: Standards Track
Expires: 23 April 2024
Authors: J. Rajamanickam, Ed. R. Gandhi, Ed.
Cisco Systems, Inc. Cisco Systems, Inc.
R. Zigler H. Song K. Kompella
Broadcom Futurewei Technologies Juniper Networks
MPLS Network Action (MNA) Sub-Stack Solution

Abstract

This document defines the MPLS Network Action (MNA) sub-stack solution for carrying Network Actions and Ancillary Data in the label stack. MPLS Network Actions can be used to influence packet forwarding decisions, carry additional OAM information in the MPLS packet, or perform user-defined operations. This document addresses the MNA requirements specified in draft-ietf-mpls-mna-requirements. This document follows the MNA framework specified in draft-ietf-mpls-mna-fwk.

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 23 April 2024.

Copyright Notice

Copyright (c) 2023 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. Conventions Used in This Document](#)
 - [2.1. Requirements Language](#)
 - [2.2. Abbreviations](#)
- [3. Overview](#)
- [4. Label Stack Entry Formats](#)
 - [4.1. LSE Format A: The MNA Sub-Stack Indicator](#)
 - [4.2. LSE Format B: The initial opcode](#)
 - [4.3. LSE Format C: Subsequent opcodes](#)
 - [4.4. LSE Format D: Additional Data](#)
- [5. The MNA Sub-Stack](#)
 - [5.1. Opcodes](#)
 - [5.2. Data](#)
 - [5.3. Scope](#)
 - [5.4. Unknown Action Handling](#)
 - [5.5. Ordering](#)
 - [5.6. Examples](#)
- [6. Special Opcodes](#)
 - [6.1. bSPL Protection](#)
 - [6.2. Flag-Based NAIs without AD](#)
 - [6.3. Extension Opcode](#)
- [7. NAS placement in the Label Stack](#)
 - [7.1. Actions when Pushing Labels](#)
- [8. Node Capability Signaling](#)
- [9. Processing the Network Action Sub-Stack](#)
 - [9.1. Encapsulating Node Responsibilities](#)
 - [9.2. Transit Node Responsibilities](#)
 - [9.3. Penultimate Node Responsibilities](#)
 - [9.4. Decapsulating Node Responsibilities](#)
- [10. Network Action Indicator Allocation Procedures](#)
- [11. Backward Compatibility](#)
- [12. Security Considerations](#)
- [13. IANA Considerations](#)
 - [13.1. MNA bSPL Label](#)
 - [13.2. MPLS Network Actions Parameters](#)
 - [13.3. Network Action Flags Without Ancillary Data](#)
 - [13.4. Network Action Opcodes](#)
- [14. Examples](#)
 - [14.1. Network Action Encoding Examples](#)
 - [14.1.1. Network Action Flags without AD](#)
 - [14.1.2. Network Action Opcode with AD](#)
 - [14.1.3. Network Action Opcode with more AD](#)

[14.2. Network Action Processing Order](#)
[14.2.1. Network Action Processing Order](#)

[15. References](#)

[15.1. Normative References](#)
[15.2. Informative References](#)

[Acknowledgments](#)

[Contributors](#)

[Authors' Addresses](#)

1. Introduction

[[RFC3032](#)] defines the encoding of the MPLS label stack, the basic structure used to define a forwarding path. Forthcoming applications require MPLS packets to perform special network actions and carry optional Ancillary Data (AD) that can affect the packet forwarding decision or trigger OAM logging, for example. Ancillary Data can be used to carry additional information, such as a network slice identifier or an entropy value for load balancing. Several MNA applications are described in [[I-D.ietf-mpls-mna-usecases](#)]. User-defined network actions allow new, local actions to be defined.

This document defines the syntax and semantics of network actions encoded within an MPLS Label Stack. Network actions can be encoded with or without Ancillary Data (AD), either in or after the label stack. In stack actions and ancillary data are contained in a Network Action Sub-Stack (NAS), which is recognized by a new base Special Purpose Label (bSPL) (value TBA). This document addresses the requirements specified in [[I-D.ietf-mpls-mna-requirements](#)]. This document follows the framework specified in [[I-D.ietf-mpls-mna-fwk](#)].

2. Conventions Used in This Document

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

2.2. Abbreviations

The terminology defined in [[I-D.ietf-mpls-mna-fwk](#)] and [[I-D.ietf-mpls-mna-requirements](#)] are used in this document.

Abbreviation	Meaning	Reference
AD	Ancillary Data	[I-D.ietf-mpls-mna-requirements]
bSPL	Base Special Purpose Label	[RFC9017]

Abbreviation	Meaning	Reference
BOS	Bottom Of Stack	[RFC3032]
HBH	Hop-By-Hop Scope	[I-D.ietf-mpls-mna-fwk]
I2E	Ingress-To-Egress Scope	[I-D.ietf-mpls-mna-fwk]
IHS	I2E, HBH, or Select Scope	This document
ISD	In-Stack Data	[I-D.ietf-mpls-mna-requirements]
LSE	Label Stack Entry	[RFC3032]
RLD	Readable Label Depth	[I-D.ietf-mpls-mna-fwk]
MNA	MPLS Network Actions	[I-D.ietf-mpls-mna-fwk]
NAI	Network Action Indicator	[I-D.ietf-mpls-mna-requirements]
NAL	Network Action Length	This document
NAS	Network Action Sub-Stack	[I-D.ietf-mpls-mna-fwk]
NASI	Network Action Sub-Stack Indicator	This document
NASL	Network Action Sub-Stack Length	This document
OAM	Operations And Management	[RFC4377]
TC	Traffic Class	[RFC5462]
TTL	Time To Live	[RFC3032]

Table 1: Abbreviations

3. Overview

The MPLS Network Action Sub-Stack (NAS) is a set of Label Stack Entries (LSEs) that appear as part of an MPLS Label Stack and serve to encode information about the network actions that should be invoked for the encapsulated packet. Multiple NASes may appear in a label stack.

Network actions and their optional Ancillary Data (AD) may be encoded as part of the NAS as a series of LSEs.

4. Label Stack Entry Formats

The NAS uses a variety of different formats of LSEs for different purposes. This section describes the syntax of the various formats while the overall structure of the NAS and the semantics of the various LSEs are described in the sections below.

4.1. LSE Format A: The MNA Sub-Stack Indicator

LSE Format A is a traditional LSE, as described in [RFC3032] and [RFC5462]. The label value is an IANA allocated value (TBA) for the

MNA bSPL label from the "Base Special-Purpose MPLS Label Values" registry to indicate the presence of an MNA Sub-Stack in the label stack.




Figure 1: LSE Format A: The MNA Sub-Stack Indicator

4.2. LSE Format B: The initial opcode

LSE Format B is used to encode the first opcode in the NAS, plus a number of other fields about the NAS.




Figure 2: LSE Format B: The initial opcode

*Opcode (7 bits) : The operation code for this LSE. See [Section 5.1](#).

*Data (13 bits) : Opcode specific data

*R (1 bit) : Reserved bit. This MUST be transmitted as zero and ignored upon receipt.

*IHS (2 bits) : The scope of the sub-stack. See [Section 5.3](#).

*S (1 bit) : The Bottom of Stack [[RFC3032](#)].

*Res (3 bits) : Reserved bits. These MUST be transmitted as zeros and ignored upon receipt.

*U (1 bit): Unknown Action Handling. See [Section 5.4](#).

*NASL (4 bits) : The Network Action Sub-Stack Length (NASL). The number of additional LSEs in the sub-stack, not including the leading Format A LSE and the Format B LSE.

NOTE: The Format A and B MUST be present while encoding Format C. The Format A, B and C MUST be present while encoding Format D.

4.3. LSE Format C: Subsequent opcodes

LSE Format C is used to encode the subsequent opcodes in the NAS.




Figure 3: LSE Format C: Subsequent opcodes

- *Opcode (7 bits) : The operation code for this LSE. See [Section 5.1](#).
- *Data (16 bits + 4 bits) : Opcode specific data
- *S (1 bit) : The Bottom of Stack [[RFC3032](#)].
- *NAL (4 bits): Network Action Length. The number of LSEs of additional data, encoded in LSE Format D ([Section 4.4](#)) following this LSE.

4.4. LSE Format D: Additional Data

LSE Format D is used to encode additional data that did not fit in the LSE with the preceding opcode.




Figure 4: LSE Format D: Additional Data

- *1 (1 bit) : The most significant bit MUST be set. This prevents legacy implementations from misinterpreting this LSE as containing a special label.
- *S (1 bit) : The Bottom of Stack [[RFC3032](#)].
- *Data (22 bits + 8 bits) : Opcode specific data

5. The MNA Sub-Stack

The MNA Sub-Stack MUST begin with a Format A LSE ([Section 4.1](#)). The label field of the LSE contains the MNA bSPL (value TBA) to indicate the presence of the MNA Sub-Stack.

The TC and TTL fields of the first LSE retain their traditional semantics, as the penultimate node on the path may copy the TTL and TC fields from the preceding LSE to the next LSE on the label stack, overwriting the TTL and TC fields of the next LSE, as specified in Section 3.5 of [[RFC3443](#)]. If the node performing this copy is not aware of MNA, this could overwrite the values in the first LSE of the MNA sub-stack.

The second LSE in a NAS MUST be a Format B LSE ([Section 4.2](#)). This LSE contains an initial opcode plus additional fields that describe the NAS.

A NAS MAY contain more Format C ([Section 4.3](#)) and Format D ([Section 4.4](#)) LSEs, up to the length encoded in the NASL field. All Format D LSEs MUST follow a Format C LSE and be included in that LSE's NAL field.

5.1. Opcodes

The opcode is a 7-bit field that indicates the semantics of its LSE. Several opcodes are assigned special semantics ([Section 6](#)), others act as Network Action Indicators and are allocated through IANA ([Section 10](#) and [Section 13.4](#)).

5.2. Data

The data field carries opcode specific data. This may be ancillary data for a network action.

To preserve backward compatibility, if a network action encodes data that will change during packet forwarding, then that data MUST be in the least significant 4 bits in the data field of a Format C LSE ([Section 4.3](#)) or the least significant 8 bits of a Format D LSE ([Section 4.4](#)). Some legacy implementations may use the label field in all LSEs when computing ECMP decisions and modifying the label field might disrupt that packet's flow.

5.3. Scope

The IHS field in the Format B LSE indicates the scope of the In-Stack NAIs encoded in the NAS. Scope defines which nodes along the MPLS path should perform the network actions found within the NAS. The specific values of the IHS field are as follows:

Bits	Scope
00	I2E
01	HBH
10	Select
11	Reserved

Table 2: IHS
Scope Values

Hop-By-Hop (HBH) - All nodes along the path MUST process the NAS.

Select - Only specific nodes along the path will perform the action.

Ingress To Egress (I2E) - The NAS MUST be processed only by the egress node.

A single NAS carries only one of the three scopes (HBH/Select/I2E). To support multiple scopes for a single packet, multiple NASes MAY be included in a single label stack.

The egress node is included in the HBH scope. This implies that the penultimate node MUST NOT remove a HBH NAS. The egress node MAY receive a NAS at the top of the label stack.

An I2E scope NAS MUST be encoded after any HBH or Select scope NASes. This makes it easier for the transit nodes to process a NAS with HBH or Select scope.

Forwarding and egress nodes should process at most a single NAS per scope. If a node is to process multiple NASes, it should process them in the order that they appear in the label stack.

5.4. Unknown Action Handling

The Unknown Action Handling (U) field in a Format B LSE ([Section 4.3](#)) is a 1-bit value that defines the action to be taken by a node that does not understand an action within the NAS. The different types of Unknown Action Handling actions are defined below.

Bit	Action
0	Skip to the next NA
1	Drop the packet

Table 3: Unknown Action
Handling

5.5. Ordering

The network actions encoded in the NAS MUST be processed as if they were processed in the order that they appear in the NAS, from the top of the NAS to the bottom. NAI encoded as flags MUST be processed as if they were processed from the most significant bit to the least significant bit. If a label stack contains multiple NASes, then they MUST be processed as if they were processed in the order that they appear in the label stack, subject to the restrictions in [Section 7](#).

5.6. Examples

A minimal NAS would have the following format, where the Label field would contain the MNA bSPL and the NASL value would be 0:




Figure 5: Example 1

A more complex NAS might have multiple opcodes and additional Ancillary Data. This example has two opcodes and two additional LSEs of AD.

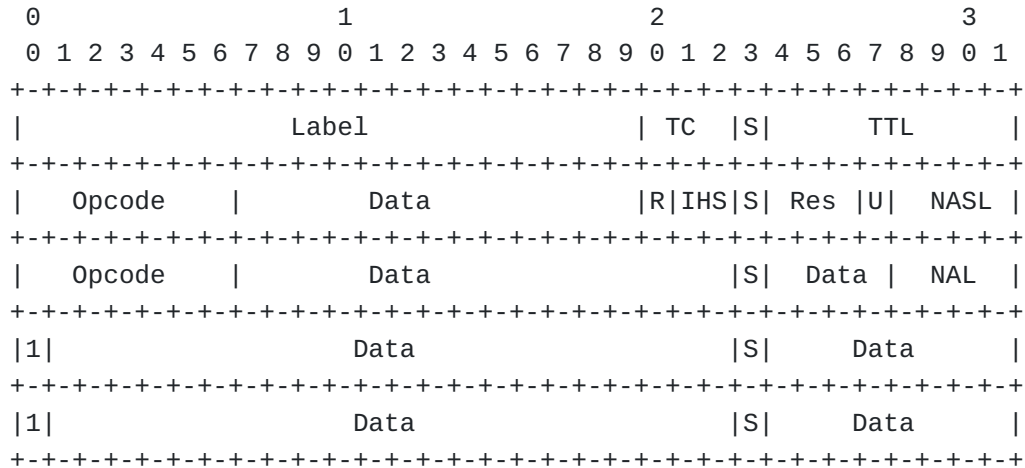


Figure 6: Example 2

In this example, the NASL field would have value 3 and the NAL field would have value 2.

6. Special Opcodes

6.1. bSPL Protection

Opcode: 0

Purpose: Legacy implementations may scan the label stack looking for bSPL values. As long as the opcode field is non-zero, an LSE cannot be misinterpreted as containing a bSPL. Opcode 0 is therefore reserved and is not used.

6.2. Flag-Based NAIs without AD

Opcode: 1

Purpose: Network actions that do not require Ancillary Data do not require an entire LSE. A single flag can be used to indicate each of these network actions.

LSE Formats: B, C, D

Data: The data field carries Network Action Indicators, which should be evaluated from the most significant bit to the least significant bit. If there are sufficient NAI, then Format D LSEs may be used to encode more flags for more network actions. Flags are allocated from the "Network Action Flags Without Ancillary Data" registry ([Section 13.3](#)). If flags need to be evaluated in a different order, multiple LSEs using this opcode may be used to specify the requested order. If this opcode is used with LSE Format B, then only 13 flags may be carried.

Scope: This opcode can be used with any scope.

This opcode MAY be used with no flags set in the data field to signify that no operation is to be performed. This can be used, for example, if the first action to be performed cannot be encoded in a Format B LSE.

6.3. Extension Opcode

Opcode: 127

Purpose: This opcode is reserved to extend the current opcode range beyond 127. Future use of this opcode is out of scope.

7. NAS placement in the Label Stack

Regardless of whether packets are being forwarded based on Segment Routing [[RFC8662](#)], LDP [[RFC5036](#)], or RSVP-TE [[RFC3209](#)], the node adding an NAS to the label stack will need to place a copy of the

NAS where it can be read by the relevant nodes. Each downstream node along the path will have Readable Label Depth (RLD) [[I-D.ietf-mpls-mna-fwk](#)] (including the LSEs of format B, C and D). If the NAS is to be processed by a particular downstream MNA capable node, then the entire NAS MUST be placed so that it is within RLD by the time the packet reaches the downstream MNA capable node and ensure the NAS MUST NOT appear at the top of the stack at any MNA incapable node on the path.

If the label stack is deep, several copies of the NAS may need to be encoded in the label stack.

For a NAS with HBH scope, every node will process the top copy of the NAS.

For a NAS with Select scope, it is processed by the node that brings it to the top of stack and then the NAS is removed from the stack. The select scoped NAS needs to be inserted after the forwarding label and needs to be inserted before the next forwarding label. It could be inserted before or after a HBH NAS.

For I2E scope, only one copy of the NAS needs to be added at the bottom of the stack.

Transit, non-penultimate nodes that pop a forwarding label and expose a copy of a NAS MUST remove it.

A node performing Penultimate Hop Popping (PHP) that pops the forwarding label with only the NAS(es) remaining on the stack MUST NOT remove the NAS(es). Instead, it forwards the packet with the NAS(es) at the top of stack to the next node.

The node that receives the NAS at the top of the label stack has to remove it.

7.1. Actions when Pushing Labels

An MNA capable node may need to push additional labels as well as push new network actions onto a received packet.

While pushing additional labels on to the label stack, the MNA capable node SHOULD verify that the entire top-most NAS with HBH scope is still within the RLD of the downstream MNA capable nodes. If required, the MNA capable node MAY create a copy of the top-most NAS with HBH scope and insert it within the RLD of the downstream MNA capable nodes on the label stack.

When an MNA capable node needs to push a new NAS with HBH scope on to a received packet that already has an NAS with HBH scope, it SHOULD copy (and merge) the network actions (including their

Ancillary Data) from the received top-most NAS with HBH scope in the new NAS with HBH scope. The new NAS MUST be placed within the RLD of the downstream MNA capable nodes. This behaviour can be based on local policy.

The new network actions added MUST NOT conflict with the network actions in the received NAS with HBH scope. The mechanism to resolve such conflicts depend on the network actions and can be based on local policy. The MNA Capable node pushing entries MUST understand any network actions which it is pushing which may result in a conflict, and MUST resolve any conflicts between new and received network actions. In the usual case of a conflict of duplicating a network action, the definition of the network action will generally give guidance on likely resolutions.

8. Node Capability Signaling

The head-end node which is adding a NAS MUST make sure that the egress node removes the NAS. The head-end node MUST make sure that the NAS can be processed by the appropriate transit and egress nodes.

*Each participating node MUST signal the network actions that it supports.

*Each participating node MUST signal its Readable Label Depth. This will allow the head-end node to place a copy of an NAS at the correct stack depth.

The above capability signaling will be added in appropriate protocols. Signaling details are outside the scope of this document.

9. Processing the Network Action Sub-Stack

This section defines the specific responsibilities for nodes along a MPLS path.

9.1. Encapsulating Node Responsibilities

The encapsulating node MAY add NASes to the label stack in accordance with its policies, the placement restrictions in [Section 7](#), and the limitations learned from [Section 8](#).

The encapsulating node MUST NOT add a NAS to the label stack if the decapsulation node does not support MNA.

If there is an existing label stack, the encapsulating node SHOULD NOT change the first 20 bits of each LSE in the label stack to avoid ECMP path change.

If the encapsulating node is also a transit node, then it MUST also respect transit node responsibilities.

The path computation needs to know the Maximum SID Depth (MSD) and Readable Label Depth (RLD) that can be imposed at the ingress node of a given SR path [[RFC8664](#)]. This ensures that the label stack depth of a computed path does not exceed the maximum number of labels (i.e., MSD) the node is capable of imposing and the maximum number of labels those could be read by the MNA processing nodes in the path. The MSD needs to include the MNA Sub-Stacks to be added.

9.2. Transit Node Responsibilities

Transit nodes SHOULD NOT change the first 20 bits in the LSEs in the label stack.

A transit node MAY change the Ancillary Data found in the least significant 8 bits of an LSE.

Transit nodes MUST process the NASes in the label stack, respecting [Section 5.5](#).

A transit node MUST respect the Unknown Action Handling value encoded in the NAS.

9.3. Penultimate Node Responsibilities

In addition to the transit node responsibilities above, the penultimate node MUST NOT remove the last copy of a HBH or I2E NAS when it is exposed after removing the forwarding (transport) label. This allows the egress node to process the NAS.

9.4. Decapsulating Node Responsibilities

The decapsulating node MUST remove any NAS it receives.

10. Network Action Indicator Allocation Procedures

This section discusses the procedures and requirements for a allocating a new opcode or flag as a network action indicator (NAI) for a network action.

A request for a new NAI MUST include the following information:

*Scope: The request MUST specify at-least one scope (I2E, HBH, Select) for the Network Action. The request MAY specify more than one scope.

*Ancillary Data: A request MUST specify the quantity, syntax, and semantics of any associated Ancillary Data. The Ancillary Data

MAY be variable length, but the length MUST be computable based on the data present in the NAS.

*Processing: The request MUST specify the detailed procedure for processing the network action.

*Interactions: The definition of the new Network Action MUST specify its interaction with other currently defined Network Action if there are any.

An assignment for an NAI MAY make requests from any combination of the "Network Action Opcodes" or "Network Action Flags Without Ancillary Data" registries. This decision should optimize for eventual encoding efficiency. If the NAI does not require any ancillary data, then a flag is preferred as only one bit is used in the encoding. If ancillary data is required, then the optimal choice may depend on how the action is likely to be combined with other actions. If the action is unlikely to be used in combination with other actions and at most 20 bits of ancillary data is required, then an opcode may be preferred as the encoding will only consume a single LSE. If the action is likely to be combined with other actions, then a flag is more likely to be optimal.

11. Backward Compatibility

This section discusses interactions between MNA capable and legacy, non-MNA capable nodes.

An MNA encapsulating node MUST ensure that the MPLS Network Action Sub-Stack indicator is not at the top of the MPLS Label Stack when the packet arrives at a non-MNA capable node. If such a packet did arrive at a non-MNA capable node, it will most likely be dropped.

Legacy nodes may scan the label stack, potentially looking for a label field containing a bSPL. To ensure that the LSE formats described herein do not appear to contain a bSPL value, the opcode value of 0 has been reserved. By ensuring that there is a non-zero value in the high order 7 bits, we are assured that the high order 20 bits cannot be misinterpreted as containing a bSPL value (0-15).

The TC and TTL fields of the Format A LSE are not re-purposed for encoding, as the penultimate node on the MPLS packet path may propagate TTL from the transport (or forwarding) label to the next label on the label stack, overwriting the TTL on the next label. If the penultimate node is a legacy node, it might perform this action, potentially corrupting other values stored in the TC and TTL fields. To protect against this, we retain the TC and TTL fields in the Format A LSE.

12. Security Considerations

The security considerations in [[RFC3032](#)] also apply to this document.

In addition, MNA creates a new dimension in security concerns:

*The actions of an encapsulating node can affect any or all of the nodes along the path. In the most common and benign situations, such as a syntactically incorrect packet, this could result in packet loss or corruption.

*The semantics of a network action are unbounded and may be insecure. A network action could be defined that made arbitrary changes to the memory of the forwarding router, which could then be used by the encapsulating node to compromise every MNA capable router in the network. The IETF needs to ensure that only secure network actions are defined.

*The MNA architecture supports locally defined network actions. For such actions, there will be limited oversight to ensure that the semantics do not create security issues. Implementors and network operators will need to ensure that locally defined network actions do not compromise the security of the network.

13. IANA Considerations

13.1. MNA bSPL Label

This document requests that IANA allocate a value (TBA) for the MNA bSPL label from the "Base Special-Purpose MPLS Label Values" registry to indicate the presence of an MNA Sub-Stack in the label stack. The description of the value should be "MPLS Network Actions". The reference should be this document.

13.2. MPLS Network Actions Parameters

This document requests that IANA create a new registry group called "MPLS Network Actions Parameters" within the "Multiprotocol Label Switching Architecture (MPLS)" registry group. The registries described below should belong to this new registry group.

13.3. Network Action Flags Without Ancillary Data

This document requests that IANA create a new registry with the name "Network Action Flags Without Ancillary Data". Registration requests should comply with [Section 10](#). The registration procedure for this registry is "IETF Review". The fields in this registry are "Bit Position" (integer), "Description" (string), and "Reference" (string).

Bit Position refers to the position relative to the most significant bit in LSE Format B or C Data fields and any subsequent Format D LSEs. Bit Position 0 is the most significant bit a LSE Format B or C Data field. Bit Position 20 is the most significant bit in the first LSE Format D Data field. There are 20 bits available in LSE Format C and 30 available in LSE Format D. There are at most 15 Format D LSEs per opcode, so there are at most $20 + 15 * 30 = 470$ bit positions. The Bit Position is an integer with value 0-469.

The initial assignments for this registry are:

Bit Position	Description	Reference
0-14	IETF Review	This document
15-16	Experimental Use	This document
17-19	Private Use	This document
20-469	IETF Review	This document

Table 4: Network Action Flags Without Ancillary Data Registry

13.4. Network Action Opcodes

This document requests that IANA create a new registry with the name "Network Action Opcodes". Registration requests should comply with [Section 10](#). The registration procedure for this registry is "IETF Review". The fields are "Opcode" (integer), "Description" (string), and "Reference" (string). Opcode is an integer 0-127.

The initial assignments for this registry are:

Opcode	Description	Reference
0	Reserved	This document
1	Flag-Based Network Action Indicators without AD	This document
2-110	IETF Review	This document
111-114	Experimental Use	This document
115-126	Private Use	This document
127	Opcode Range Extension Beyond 127	This document

Table 5: Network Action Opcodes Registry

14. Examples

14.1. Network Action Encoding Examples

14.1.1. Network Action Flags without AD




Figure 7: NAS with Network Action Flags

This is an example of an NAS with Flag-Based NAIs without Ancillary Data.

Details:

Opcode=2: This opcode to indicates that the LSE carries Flag-Based NAIs without AD.

Data: The data field carries the Flag-Based NAIs.

S: This is the bottom of stack bit. Set if and only if this LSE is the bottom of the stack.

U: Action to be taken if one of the NAIs are not recognized by the processing node.

NASL: The NASL field is set to "0", as there are no additional LSEs.

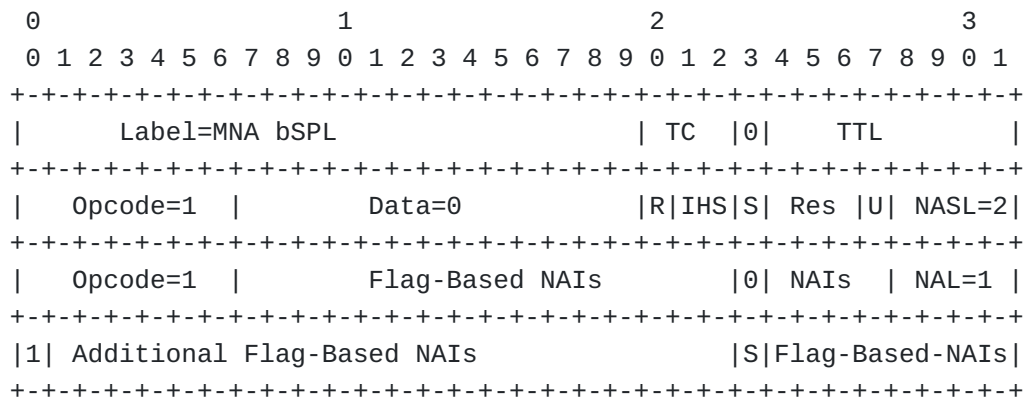


Figure 8: Network Action Flags without AD using LSE Format D

In this example, the NAS contains a Format B LSE with no flags set, indicating no operation. The next LSE uses Format C, but the Network Action Flag is not in a bit position contained within the Format C LSE, so a single Format D LSE has been added to the NAS to carry the flag.

NAL is set to "1" to indicate that Flag-Based NAIs are also encoded in the next LSE.

NASL is set to "2" to indicate that 2 additional LSEs are used.

14.1.2. Network Action Opcode with AD




Figure 9: Network action opcode with Ancillary Data

In this example, the NAS is carrying only one Network Action that requires 13 bits of Ancillary Data.

Details on the Second LSE

Opcode=8: A network action allocated outside of this document.

Data: The data field contains 13 bits of ancillary data.

14.1.3. Network Action Opcode with more AD

A network action may require more Ancillary Data than can fit in a single LSE. In this example, a Format D LSE is added to carry additional Ancillary Data.

```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|          Label=MNA bSPL          | TC |0|      TTL      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode=1 |          Data=0          |R|IHS|0| Res |U| NASL=2|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode=9 |          Ancillary Data          |0| AD  | NAL=1 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|1|          Ancillary Data          |S|Ancillary Data |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 10: Network Action With Additional Ancillary Data

In this example, opcode 9 requires more than one LSE's worth of Ancillary Data, so a Format D LSE is added.

Details on the third LSE:

Opcode=9: An opcode allocated outside of this document

Ancillary Data: Most significant bits of Ancillary data

AD: 4 bits of additional Ancillary Data

Details on the fourth LSE:

Ancillary Data: 22 bits of additional Ancillary data.

Ancillary Data: 8 bits of additional Ancillary Data.

14.2. Network Action Processing Order

The semantics of a network action can vary widely and the results of processing one network action may affect the processing of a subsequent network action. See [Section 5.5](#).

14.2.1. Network Action Processing Order

```

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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Label=MNA bSPL           | TC |S|   TTL           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode=8   |   Ancillary Data   |R|IHS|S|Res|U|1| NASL=2|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode=7   |   Ancillary Data7   |S| AD7 | NAL=0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode=1   |   Flag-Based NAIs   |S| NAI | NAL=0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 11: In-stack NA processing order

In this example, opcode 8 is processed first, then opcode 7, and then the network action flags are processed from most significant to least significant.

In a different case, some Flag-Based NAIs may need to be processed before opcode 7 and some Flag-Based NAIs may need to be processed after Opcode 7. This can be done by causing some NAIs to appear earlier in the 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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|           Label=MNA bSPL           | TC |S|   TTL           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode=8   |   Ancillary Data   |R|IHS|S|Res|U|1| NASL=3|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode=1   |       0x01       |S| NAI | NAL=0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode=7   |   Ancillary Data7   |S| AD7 | NAL=0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Opcode=1   |       0x02       |S| NAI | NAL=0 |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 12: Interleaving network actions

In the above example, opcode 8 is processed first, then Flag-Based NAI 0x01 is processed before opcode 7, and finally NAI 0x02 is processed.

15. References

15.1. Normative References

[I-D.ietf-mp1s-mna-fwk]

Andersson, L., Bryant, S., Bocci, M., and T. Li, "MPLS Network Actions Framework", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-fwk-05, 19 October 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-fwk-05>>.

- [**I-D.ietf-mpls-mna-requirements**] Bocci, M., Bryant, S., and J. Drake, "Requirements for MPLS Network Actions", Work in Progress, Internet-Draft, draft-ietf-mpls-mna-requirements-07, 18 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-requirements-07>>.
- [**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>>.
- [**RFC3209**] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [**RFC3443**] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, DOI 10.17487/RFC3443, January 2003, <<https://www.rfc-editor.org/info/rfc3443>>.
- [**RFC5036**] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.
- [**RFC5462**] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, DOI 10.17487/RFC5462, February 2009, <<https://www.rfc-editor.org/info/rfc5462>>.
- [**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>>.
- [**RFC8662**] Kini, S., Kompella, K., Sivabalan, S., Litkowski, S., Shakir, R., and J. Tantsura, "Entropy Label for Source Packet Routing in Networking (SPRING) Tunnels", RFC 8662,

DOI 10.17487/RFC8662, December 2019, <<https://www.rfc-editor.org/info/rfc8662>>.

[RFC9017] Andersson, L., Kompella, K., and A. Farrel, "Special-Purpose Label Terminology", RFC 9017, DOI 10.17487/RFC9017, April 2021, <<https://www.rfc-editor.org/info/rfc9017>>.

15.2. Informative References

[RFC4377] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and S. Matsushima, "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, DOI 10.17487/RFC4377, February 2006, <<https://www.rfc-editor.org/info/rfc4377>>.

[RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.

[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-03, 15 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-usecases-03>>.

Acknowledgments

The authors of this document would like to thank the MPLS Working Group Open Design Team for the discussions and comments on this document. The authors would also like to thank Amanda Baber for reviewing the IANA Considerations and providing many useful suggestions. The authors would like to thank Loa Andersson, Stewart Bryant, Greg Mirsky and Joel M. Halpern for reviewing this document and providing many useful suggestions.

Contributors

The following people have substantially contributed to this document:

Jisu Bhattacharya
Cisco Systems, Inc.
Email: jisu@cisco.com

Bruno Decraene
Orange
Email: bruno.decraene@orange.com

Weiqiang Cheng
China Mobile
Email: chengweiqiang@chinamobile.com

Xiao Min
ZTE Corp.
Email: xiao.min2@zte.com.cn

Luay Jalil
Verizon
Email: luay.jalil@verizon.com

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China
Email: jie.dong@huawei.com

Tianran Zhou
Huawei Technologies
China
Email: zhoutianran@huawei.com

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

Sami Boutros
Ciena
Email: sboutros@ciena.com

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

John Drake
Juniper Networks
United States
Email: jdrake@juniper.net

Figure 13

Authors' Addresses

Jaganbabu Rajamanickam (editor)
Cisco Systems, Inc.
Canada

Email: jrajaman@cisco.com

Rakesh Gandhi (editor)
Cisco Systems, Inc.
Canada

Email: rgandhi@cisco.com

Royi Zigler
Broadcom

Email: royi.zigler@broadcom.com

Haoyu Song
Futurewei Technologies

Email: haoyu.song@futurewei.com

Kireeti Kompella
Juniper Networks
United States

Email: kireeti.ietf@gmail.com