
Stream: Internet Engineering Task Force (IETF)
RFC: [9789](#)
Category: Informational
Published: May 2025
ISSN: 2070-1721
Authors:
L. Andersson S. Bryant M. Bocci T. Li
Huawei Technologies University of Surrey 5GIC Nokia Juniper Networks

RFC 9789

MPLS Network Action (MNA) Framework

Abstract

This document describes an architectural framework for MPLS Network Action (MNA) technologies. MNA technologies are used to indicate actions that impact the forwarding or other processing (such as monitoring) of the packet along the Label Switched Path (LSP) of the packet and to transfer any additional data needed for these actions.

This document provides the foundation for the development of a common set of network actions and information elements supporting additional operational models and capabilities of MPLS networks.

Status of This Memo

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

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9789>.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions

with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Language	4
1.2. Normative Definitions	4
1.3. Abbreviations	4
2. Structure	5
2.1. Scopes	7
2.2. Partial Processing	8
2.3. Signaling	8
2.3.1. Readable Label Depth	8
2.4. State	9
3. Encoding	9
3.1. The MNA Label	10
3.1.1. Existing Base SPL	10
3.1.2. New Base SPL	10
3.1.3. New Extended SPL	10
3.1.4. User-Defined Label	10
3.2. TC and TTL	10
3.2.1. TC and TTL Retained	10
3.2.2. TC and TTL Repurposed	11
3.3. Length of the NAS	11
3.3.1. Last/Continuation Bits	12
3.3.2. Length Field	12
3.4. Encoding of Scopes	12
3.5. Encoding a Network Action	12
3.5.1. Bit Catalogs	12

3.5.2. Operation Codes	13
3.6. Encoding of Post-Stack Data	13
3.6.1. First Nibble Considerations	13
4. Semantics	14
5. Definition of a Network Action	14
6. Management Considerations	15
7. Security Considerations	15
8. IANA Considerations	16
9. References	16
9.1. Normative References	16
9.2. Informative References	17
Acknowledgements	19
Authors' Addresses	19

1. Introduction

This document describes an architectural framework for MPLS Network Action (MNA) technologies. MNA technologies are used to indicate actions for Label Switched Paths (LSPs) and/or MPLS packets and to transfer data needed for these actions.

This document provides the foundation for the development of a common set of network actions and information elements supporting additional operational models and capabilities of MPLS networks. MNA solutions derived from this framework are intended to address the requirements found in [RFC9613]. In addition, MNA may support actions that overlap existing MPLS functionality. This may be beneficial for numerous reasons, such as making it more efficient to combine existing functionality and new functions in the same MPLS packet.

MPLS forwarding actions are instructions to MPLS routers to apply additional actions when forwarding a packet. These might include load-balancing a packet given its entropy, indicating whether or not to perform Fast Reroute on a failure, and indicating whether or not a packet has metadata relevant to the forwarding actions along the path.

This document generalizes the concept of MPLS "forwarding actions" to "network actions" that include any action that an MPLS router is requested to take on the packet. Network actions include any MPLS forwarding actions but may also include other operations (such as security functions, Operations, Administration, and Maintenance (OAM) procedures, etc.) that are not

directly related to forwarding of the packet. MPLS network actions are always triggered by an MNA packet but may have implications for subsequent traffic, including non-MNA packets, as discussed in [Section 2.4](#).

MNA technologies may redefine the semantics of the Label, Traffic Class (TC), and Time to Live (TTL) fields in an MPLS Label Stack Entry (LSE) within a Network Action Sub-Stack (NAS).

1.1. 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.

Although this is an Informational document, these conventions are applied to achieve clarity in the requirements that are presented.

1.2. Normative Definitions

This document adopts the definitions of the following terms and abbreviations from [[RFC9613](#)] as normative: "Network Action", "Network Action Indicator (NAI)", "Ancillary Data (AD)", and "Scope".

In addition, this document defines the following terms:

Network Action Sub-Stack (NAS): A set of related, contiguous LSEs in the MPLS label stack for carrying information related to network actions. The Label, TC, and TTL values in the LSEs in the NAS may be redefined, but the meaning of the S bit is unchanged.

Network Action Sub-Stack Indicator (NSI): The first LSE in the NAS contains a special-purpose label that indicates the start of the NAS.

1.3. Abbreviations

Abbreviation	Meaning	Reference
AD	Ancillary Data	[RFC9613]
BIER	Bit Index Explicit Replication	[RFC8279]
BoS	Bottom of Stack	[RFC6790]
bSPL	Base Special-Purpose Label	[RFC9017]
ECMP	Equal-Cost Multipath	[RFC9522]
EL	Entropy Label	[RFC6790]

Abbreviation	Meaning	Reference
ERLD	Entropy Readable Label Depth	[RFC8662]
eSPL	Extended Special-Purpose Label	[RFC9017]
HbH	Hop by Hop	In the MNA context, this document.
I2E	Ingress to Egress	In the MNA context, this document.
IGP	Interior Gateway Protocol	
ISD	In-Stack Data	[RFC9613]
LSE	Label Stack Entry	[RFC3032]
LSP	Label Switched Path	
MNA	MPLS Network Action	[RFC9613]
MSD	Maximum SID Depth	[RFC8491]
NAI	Network Action Indicator	[RFC9613]
NAS	Network Action Sub-Stack	This document
NSI	Network Action Sub-Stack Indicator	This document
PSD	Post-Stack Data	[RFC9613] and Section 3.6
RLD	Readable Label Depth	This document
SID	Segment Identifier	[RFC8402]
SPL	Special-Purpose Label	[RFC9017]
TC	Traffic Class	
TTL	Time to Live	

Table 1: Abbreviations

2. Structure

An MNA solution specifies one or more network actions to apply to an MPLS packet. These network actions and their ancillary data may be carried in sub-stacks within the MPLS label stack and/or post-stack data. A solution must specify where the network action sub-stacks occur in the label stack, if and how frequently they should be replicated within the label stack, and how the network action sub-stack and post-stack data are encoded.

It seems highly likely that some ancillary data will be needed at many points along an LSP. Replication of ancillary data throughout the label stack would be highly inefficient, as would a full rewrite of the label stack at each hop; thus, MNA allows encoding of network actions and ancillary data deeper in the label stack, requiring implementations to look past the first LSE. Processing of the label stack past the top-of-stack LSE was first introduced with the Entropy Label (EL) [RFC6790].

A network action sub-stack contains:

- **Network Action Sub-Stack Indicator (NSI):** The first LSE in the NAS contains a special-purpose label, called the MNA label, which is used to indicate the start of a network action sub-stack.
- **Network Action Indicators (NAIs):** Optionally, a set of indicators that describes the set of network actions. If the set of indicators is not in the sub-stack, a solution could encode them in post-stack data. A network action is said to be present if there is an indicator in the packet that invokes the action.
- **In-Stack Data (ISD):** A set of zero or more LSEs that carry ancillary data for the network actions that are present. Network action indicators are not considered ancillary data.

Each network action present in the network action sub-stack may have zero or more LSEs of in-stack data. The ordering of the in-stack data LSEs corresponds to the ordering of the network action indicators. The encoding of the in-stack data, if any, for a network action must be specified in the document that defines the network action. In-stack data may be referenced by multiple network actions.

As an example, in-stack data might look like the following label stack with an embedded NAS:




Figure 1: A Label Stack with an Embedded Network Action Sub-Stack

Certain network actions may also specify that data is carried after the label stack. This is called post-stack data. The encoding of the post-stack data, if any, for a network action must be specified in the document that defines the network action. If multiple network actions are present and have post-stack data, the ordering of their post-stack data corresponds to the ordering of the network action indicators.

As an example, a packet containing post-stack data might contain a label stack followed by post-stack data, followed by the payload:

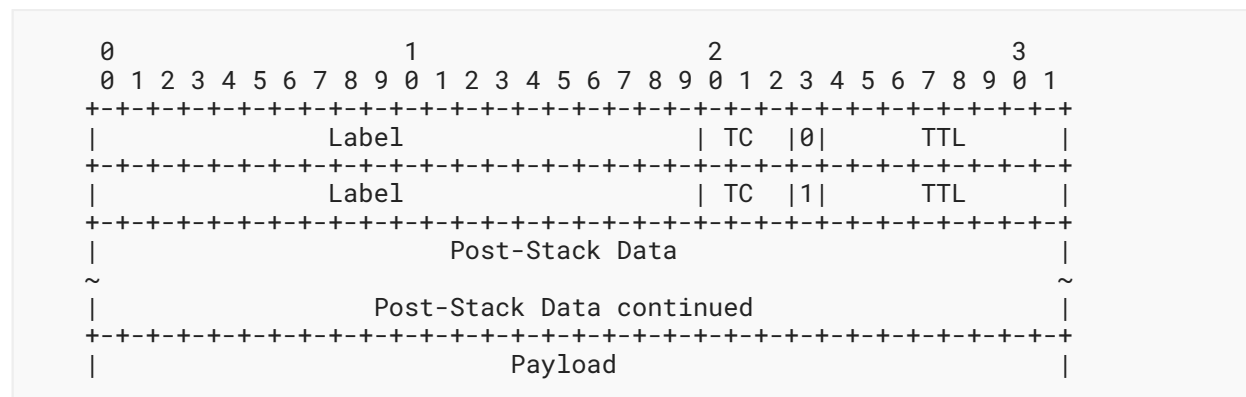


Figure 2: A Label Stack Followed by Post-Stack Data

A solution must specify the order for network actions to be applied to the packet for the actions to have consistent semantics. Since there are many possible orderings, especially with bit catalogs ([Section 3.5.1](#)), the solution must provide an unambiguous specification. The precise semantics of an action are dependent on the contents of the packet, including any ancillary data, and the state of the router.

This document assumes that the MPLS WG will select a single solution for the encoding of ISD and not more than one solution for the encoding of PSD.

2.1. Scopes

A network action may need to be processed by every node along the path or some subset of the nodes along its path. Some of the scopes that an action may have are:

- Hop by Hop (HbH): Every node along the path will perform the action.
- Ingress to Egress (I2E): Only the last node on the path will perform the action.
- Select: Only specific nodes along the path will perform the action.

If a solution supports the select scope, it must describe how it specifies the set of nodes to perform the actions.

This framework does not place any constraints on the scope of, or the ancillary data for, a network action. Any network action may appear in any scope or combination of scopes, may have no ancillary data, and may require in-stack data and/or post-stack data. Some combinations may be suboptimal, but this framework does not restrict the combinations in an MNA solution. A specific MNA solution may define such constraints.

2.2. Partial Processing

As described in [\[RFC3031\]](#), legacy devices that do not recognize the MNA label will discard the packet if the top label is the MNA label.

Devices that do recognize the MNA label might not implement all of the network actions that are present. A solution must specify how unrecognized network actions that are present should be handled.

One alternative is that an implementation should stop processing network actions when it encounters an unrecognized network action. Subsequent present network actions would not be applied. The result is dependent on the solution's order of operations.

Another alternative is that an implementation should drop any packet that contains any unrecognized present network actions.

A third alternative is that an implementation should perform all recognized present network actions but ignore all unrecognized present network actions.

Other alternatives may also be possible. The solution should specify the alternative adopted.

In some solutions, an indication may be provided in the packet or in the action as to how the forwarder should proceed if it does not recognize the action. Where an action needs to be processed at every hop, it is recommended that care be taken not to construct an LSP that traverses nodes that do not support that action. It is recognized that, in some circumstances, it may not be possible to construct an LSP that avoids such nodes, such as when a network is reconverging following a failure or when IP Fast Reroute (IPFRR) [\[RFC5714\]](#) is taking place.

2.3. Signaling

A node that wishes to make use of MNA and apply network actions to a packet must understand the nodes that the packet will transit, whether or not the nodes support MNA, and the network actions that are to be invoked. These capabilities are presumed to be signaled by protocols that are out of scope for this document and are presumed to have per-network-action granularity. If a solution requires alternate signaling, it must specify that explicitly.

2.3.1. Readable Label Depth

Readable Label Depth (RLD) is defined as the number of LSEs, starting from the top of the stack, that a router can read in an incoming MPLS packet with no performance impact. [\[RFC8662\]](#) introduced Entropy Readable Label Depth (ERLD). Readable Label Depth is the same concept, but it is generalized and not specifically associated with the Entropy Label (EL) or MNA.

ERLD is not redundant with RLD because ERLD specifies a value of zero if a system does not support the Entropy Label. Since a system could reasonably support MNA or other MPLS functions and needs to advertise an RLD value but not support the Entropy Label, another advertised value is required.

A node that pushes a NAS onto the label stack is responsible for ensuring that all nodes that are expected to process the NAS will have the entire NAS within their RLD. A node **SHOULD** use signaling (e.g., the signaling described in [RFC9088] and [RFC9089]) to determine this. An exception might be, for example, when the node has out-of-band knowledge that all nodes along the path do not have RLD limitations and thus could avoid the unnecessary overhead of using signaling.

Per [RFC8662], a node that does not support EL will advertise a value of zero for its ERLD, so advertising ERLD alone does not suffice in all cases. A node **MAY** advertise both ERLD and RLD, and it **SHOULD** do so if its ERLD and RLD values are different. Again, if a node has out-of-band knowledge that all nodes do not have RLD limitations, then signaling can be avoided. If a node's ERLD and RLD values are the same, it **MAY** only advertise ERLD for efficiency reasons. If a node supports MNA but does not support EL, then it **SHOULD** advertise RLD unless it has out-of-band knowledge that no nodes in the domain have RLD restrictions.

RLD is advertised by an IGP MSD-Type value of 3 and **MAY** be advertised as a Node MSD, Link MSD, or both.

An MNA node **MUST** use the RLD determined by selecting the first advertised non-zero value from the following:

- The RLD advertised for the link
- The RLD advertised for the node
- The non-zero ERLD for the node

A node's RLD is a function of its hardware capabilities and is not expected to depend on the specifics of the MNA solution.

2.4. State

A network action can affect the state stored in the network. This implies that a packet may affect how subsequent packets are handled. In particular, one packet may affect subsequent packets in the same LSP.

3. Encoding

Several possible ways to encode NAIs have been proposed. This section summarizes the proposals and some considerations for the various alternatives.

When network actions are carried in the MPLS label stack, then regardless of their type, they are represented by a set of LSEs termed a Network Action Sub-Stack (NAS). A NAS consists of a special-purpose label, optionally followed by LSEs that specify which network actions are to be

performed on the packet and the in-stack ancillary data for each indicated network action. Different network actions may be placed together in one NAS or may be carried in different sub-stacks.

[RFC9613] requires that a solution not add unnecessary LSEs to the sub-stack (see requirement 9 in Section 3.1 of [RFC9613]). Accordingly, solutions should also make efficient use of the bits within the sub-stack (except the S-bit), as inefficient use of the bits could result in the addition of unnecessary LSEs.

3.1. The MNA Label

The first LSE in a network action sub-stack contains a special-purpose label that indicates a network action sub-stack. A solution has several choices for this special-purpose label.

3.1.1. Existing Base SPL

A solution may reuse an existing Base SPL (bSPL). If it elects to do so, it must explain how the usage is backward compatible, including in the case where there is ISD.

If an existing inactive bSPL is selected that will not be backward compatible, then it must first be retired per [RFC7274] and then reallocated.

3.1.2. New Base SPL

A solution may select a new bSPL.

3.1.3. New Extended SPL

A solution may select a new Extended SPL (eSPL). If it elects to do so, it must address the requirement for the minimal number of LSEs.

3.1.4. User-Defined Label

A solution may allow the network operator to define the label that indicates the network action sub-stack. This creates management overhead for the network operator to coordinate the use of this label across all nodes on the path using management or signaling protocols. The user-defined label could be network-wide or LSP-specific. If a solution elects to use a user-defined label, the solution should justify this overhead.

3.2. TC and TTL

In the first LSE of the network action sub-stack, only the 20 bits of the Label value and the Bottom of Stack bit are used by the NSI; the TC field (3 bits) and the TTL (8 bits) are not used. This could leave 11 bits that could be used for MNA purposes.

3.2.1. TC and TTL Retained

If the solution elects to retain the TC and TTL fields, then the first LSE of the network action sub-stack would appear as described in [RFC3032]:




Figure 3: A Label Stack Entry with TC and TTL Retained

Label: Label value, 20 bits

TC: Traffic Class, 3 bits

S: Bottom of Stack, 1 bit

TTL: Time To Live, 8 bits

Further LSEs would be needed to encode NAIs. If a solution elects to retain the TC and TTL fields, it must address the requirement for the minimal number of LSEs.

3.2.2. TC and TTL Repurposed

If the solution elects to reuse the TC and TTL fields, then the first LSE of the network action sub-stack would appear as follows:




Figure 4: A Label Stack Entry with TC and TTL Repurposed

Label: Label value, 20 bits

x: Bit available for use in solution definition

S: Bottom of Stack, 1 bit

The solution may use more LSEs to contain NAIs. If a solution elects to use more LSEs, it must address the requirement for the minimal number of LSEs.

3.3. Length of the NAS

A solution must have a mechanism (such as an indication of the length of the NAS) to enable an implementation to find the end of the NAS. This must be easily processed even by implementations that do not understand the full contents of the NAS. Two options are described below; other solutions may be possible.

3.3.1. Last/Continuation Bits

A solution may use a bit per LSE to indicate whether or not the NAS continues into the next LSE. The bit may indicate continuation by being set or by being clear. The overhead of this approach is one bit per LSE and has the advantage that it can effectively encode an arbitrarily sized NAS. This approach is efficient if the NAS is small.

3.3.2. Length Field

A solution may opt to have a fixed-size Length field at a fixed location within the NAS. The fixed size of the Length field may not be large enough to support all possible NAS contents. This approach may be more efficient if the NAS is long, but not longer than can be described by the Length field.

One hardware designer recommends a Length field as this minimizes branching in the logic.

3.4. Encoding of Scopes

A solution may choose to explicitly encode the scope of each action contained in a network action sub-stack. For example, a NAS might contain Action A (HbH), Action B (HbH), and Action C (HbH). A solution may alternately choose to have the scope encoded implicitly, based on the actions present in the network action sub-stack. For example, a NAS might contain the following actions with HbH scope: A, B, and C. This choice may have performance implications as an implementation might have to parse the network actions that are present in a network action sub-stack only to discover that there are no actions for it to perform.

For example, suppose that a NAS is embedded in a label stack at a depth of six LSEs and the NAS contains three actions, each with Select scope. These actions are not applicable at the current node and should be ignored. If the scope is encoded explicitly with each action, then an implementation must parse each action. However, if the scope is encoded as part of the NAS, then an implementation only needs to parse the start of the NAS and not individual actions.

Solutions need to consider the order of scoped NAIs and their associated AD within individual sub-stacks and the order of per-scope sub-stacks, so that network actions and the AD can be readily found and not be processed by nodes that are not required to handle those actions.

3.5. Encoding a Network Action

Two options for encoding NAIs are described below; other solutions may be possible. Any solution should allow the encoding of an arbitrary number of NAIs.

3.5.1. Bit Catalogs

A solution may opt to encode the set of network actions as a list of bits, sometimes known as a catalog. The solution must provide a mechanism to determine how many LSEs are devoted to the catalog when the NAIs are carried in-stack. A set bit in the catalog would indicate that the corresponding network action is present.

Catalogs are efficient if the number of present network actions is relatively high and if the size of the necessary catalog is small. For example, if the first 16 actions are all present, a catalog can encode this in 16 bits. However, if the number of possible actions is large, then a catalog can become inefficient. Selecting only one action that is the 256th action would require a catalog of 256 bits, which would require more than one LSE when the NAIs are carried in-stack.

A solution may include a bit-remapping mechanism so that a given domain may optimize for its commonly used actions.

3.5.2. Operation Codes

A solution may opt to encode the set of present network actions as a list of operation codes (opcodes). Each opcode is a fixed number of bits. The size of the opcode bounds the number of network actions that the solution can support.

Opcodes are efficient if there are only one or two active network actions. For example, if an opcode is 8 bits, then two active network actions could be encoded in 16 bits. However, if 16 actions are required, then opcodes would consume 128 bits. Opcodes are efficient at encoding a large number of possible actions. If only the 256th action is to be selected, that still requires 8 bits.

3.6. Encoding of Post-Stack Data

A solution may carry NAI and AD as PSD. For ease of parsing, all AD should be co-located with its NAI.

If there are multiple instances of post-stack data, they should occur in the same order as their relevant network action sub-stacks and then in the same order as their relevant network actions occur within the network action sub-stacks.

3.6.1. First Nibble Considerations

The first nibble after the label stack has been used to convey information in certain cases [RFC4385]. A consolidated view of the uses of the first nibble is provided in [RFC9790].

For example, in [RFC4928], this nibble is investigated to find out if it has the value "4" or "6". If it does not, it is assumed that the packet payload is not IPv4 or IPv6, and Equal-Cost Multipath (ECMP) is not performed.

It should be noted that this is an inexact method. For example, an Ethernet pseudowire without a control word might have "4" or "6" in the first nibble and thus will be subject to ECMP forwarding.

Nevertheless, the method is implemented and deployed; it is used today and will be for the foreseeable future.

The use of the first nibble for Bit Index Explicit Replication (BIER) is specified in [RFC8296]. BIER sets the first nibble to 5. The same is true for a BIER payload as for any use of the first nibble: it is not possible to conclude that the payload is BIER even if the first nibble is set to 5 because an

Ethernet pseudowire without a control word might begin with a 5. However, the BIER approach meets the design goal of [\[RFC8296\]](#) to determine that the payload is IPv4, IPv6, or the header of a pseudowire packet with a control word, rather than being a payload belonging to a BIER or some other type of packet.

[\[RFC4385\]](#) allocates 0b0000 for the pseudowire control word and 0b0001 as the control word for the pseudowire Associated Channel Header (ACH).

A PSD solution should specify the contents of the first nibble, the actions to be taken for the value, and the interaction with post-stack data used concurrently by other MPLS applications.

4. Semantics

For MNA to be consistent across implementations and predictable in operational environments, its semantics need to be entirely predictable. An MNA solution **MUST** specify a deterministic order for processing each of the network actions in a packet. Each network action must specify how it interacts with all other previously defined network actions. Private network actions are network actions that are not publicly documented. Private network actions **MUST** be included in the ordering of network actions, but the interactions of private actions with other actions are outside of the scope of this document.

5. Definition of a Network Action

A document defining a network action must contain the following:

Name: The name of the network action.

Network Action Indicator: The bit position or opcode that indicates that the network action is active.

Scope: Description of which nodes should perform the network action as described in [Section 2.1](#).

State: Indication of whether the network action can modify state in the network and, if so, the state that may be modified and its side effects.

Required/Optional: Indication of whether a node is required to perform the network action.

In-Stack Data: The number of LSEs of in-stack data, if any, and the encoding of the in-stack data. If the in-stack data is of a variable length, then the solution must specify how an implementation can determine the length without implementing the network action.

Post-Stack Data: The encoding of post-stack data, if any. If the post-stack data is of a variable length, then the solution must specify how an implementation can determine the length without implementing the network action.

A solution should create an IANA registry for network actions.

6. Management Considerations

Network operators will need to be cognizant of which network actions are supported by which nodes and will need to ensure that this is signaled. Some solutions may require network-wide configuration to synchronize the use of the labels that indicate the start of a NAS. Solution documents must clearly state what management considerations apply to the solutions they are describing. Solution documents must describe mechanisms for performing network diagnostics in the presence of MNAs.

7. Security Considerations

An analysis of the security of MPLS systems is provided in [\[RFC5920\]](#), which also notes that the MPLS forwarding plane has no built-in security mechanisms.

Central to the security of MPLS networks is operational security of the network, something that operators of MPLS networks are well versed in. The deployment of link-level security (e.g., [\[MACsec\]](#)) prevents link traffic observation covertly acquiring the label stack for an attack. This is particularly important in the case of a network deploying MNA, because the MNA information may be sensitive. Thus, the confidentiality and authentication achieved through the use of link-level security is particularly advantageous.

Some additional proposals to add encryption to the MPLS forwarding plane have been suggested [\[MPLS-OPP-SEC\]](#), but no mechanisms have been agreed upon at the time of publication of this document. [\[MPLS-OPP-SEC\]](#) offers hop-by-hop security that encrypts the label stack and is functionally equivalent to that provided by [\[MACsec\]](#). Alternatively, it also offers end-to-end encryption of the MPLS payload with no cryptographic integrity protection of the MPLS label stack.

Particular care is needed when introducing any end-to-end security mechanism to allow an in-stack MNA solution that needs to employ on-path modification of the MNA data or where post-stack MNA data needs to be examined on-path.

A cornerstone of MPLS security is to protect the network from processing MPLS labels that originated outside the network.

Operators have considerable experience in excluding MPLS-encoded packets at the network boundaries, for example, by excluding all MPLS packets and all packets that are revealed to be carrying an MPLS packet as the payload of IP tunnels. Where such packets are accepted into an MPLS network from an untrusted third party, non-MPLS packets are immediately encapsulated in an MPLS label stack specified by the MPLS network operator, and MPLS packets have additional label stack entries imported as specified by the MPLS network operator. Thus, it is difficult for an attacker to pass an MPLS-encoded packet into a network or to present any instructions to the network forwarding system.

Within a single well-managed domain, an adjacent domain may be considered to be trusted provided that it is sufficiently shielded from third-party traffic ingress and third-party traffic observation. In such a situation, no new security vulnerabilities are introduced by MNA.

In some inter-domain applications (including carrier's carrier) where a first network's MPLS traffic is encapsulated directly over a second MPLS network by simply pushing additional MPLS LSEs, the contents of the first network's payload and label stack may be visible to the forwarders in the second network. Historically, this has been benign and indeed useful for ECMP. However, if the first network's traffic has MNA information, this may be exposed to MNA-capable forwarders and cause unpredictable behavior or modification of the customer MPLS label stack or MPLS payload. This is an increased vulnerability introduced by MNA that **SHOULD** be addressed in any MNA solution.

Several mitigations are available to an operator:

- a. Reject all incoming packets containing MNA information that do not come from a trusted network. Note that it may be acceptable to accept and process MNA information from a trusted network.
- b. Fully encapsulate the inbound packet in a new additional MPLS label stack such that the forwarder finds a Bottom of Stack (BoS) bit imposed by the carrier network and only finds MNA information added by the carrier network.

A mitigation that we reject as unsafe is having the ingress Label Switching Router (LSR) push sufficient additional labels such that any MNA information received in packets entering the network from a third-party network is made inaccessible due to it being below the RLD. This is unsafe in the presence of an overly conservative RLD value and can result in the third-party MNA information becoming visible to and acted on by an MNA forwarder in the carrier network.

8. IANA Considerations

IANA has allocated the following code point in the "IGP MSD-Types" registry [MSD] within the "Interior Gateway Protocol (IGP) Parameters" registry group:

Value	Name	Data Plane	Reference
3	Readable Label Depth	MPLS	RFC 9789

Table 2: New IGP MSD-Type

9. References

9.1. 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/info/rfc2119>>.

- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [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>>.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, DOI 10.17487/RFC4385, February 2006, <<https://www.rfc-editor.org/info/rfc4385>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<https://www.rfc-editor.org/info/rfc5920>>.
- [RFC7274] Kompella, K., Andersson, L., and A. Farrel, "Allocating and Retiring Special-Purpose MPLS Labels", RFC 7274, DOI 10.17487/RFC7274, June 2014, <<https://www.rfc-editor.org/info/rfc7274>>.
- [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>>.
- [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>>.
- [RFC9613] Bocci, M., Ed., Bryant, S., and J. Drake, "Requirements for Solutions that Support MPLS Network Actions (MNAs)", RFC 9613, DOI 10.17487/RFC9613, August 2024, <<https://www.rfc-editor.org/info/rfc9613>>.

9.2. Informative References

- [MACsec] IEEE, "IEEE Standard for Local and metropolitan area networks-Media Access Control (MAC) Security", IEEE Std 802.1AE-2018, DOI 10.1109/ieeestd.2018.8585421, 26 December 2018, <<https://ieeexplore.ieee.org/document/8585421>>.
- [MPLS-OPP-SEC] Farrel, A. and S. Farrell, "Opportunistic Security in MPLS Networks", Work in Progress, Internet-Draft, draft-ietf-mpls-opportunistic-encrypt-03, 28 March 2017, <<https://datatracker.ietf.org/doc/html/draft-ietf-mpls-opportunistic-encrypt-03>>.
- [MSD] IANA, "IGP MSD-Types", <<https://www.iana.org/assignments/igp-parameters/>>.
- [RFC4928] Swallow, G., Bryant, S., and L. Andersson, "Avoiding Equal Cost Multipath Treatment in MPLS Networks", BCP 128, RFC 4928, DOI 10.17487/RFC4928, June 2007, <<https://www.rfc-editor.org/info/rfc4928>>.

-
- [RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, DOI 10.17487/RFC5714, January 2010, <<https://www.rfc-editor.org/info/rfc5714>>.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/info/rfc6790>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", RFC 8279, DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", RFC 8296, DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8491] Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling Maximum SID Depth (MSD) Using IS-IS", RFC 8491, DOI 10.17487/RFC8491, November 2018, <<https://www.rfc-editor.org/info/rfc8491>>.
- [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>>.
- [RFC9088] Xu, X., Kini, S., Psenak, P., Filsfils, C., Litkowski, S., and M. Bocci, "Signaling Entropy Label Capability and Entropy Readable Label Depth Using IS-IS", RFC 9088, DOI 10.17487/RFC9088, August 2021, <<https://www.rfc-editor.org/info/rfc9088>>.
- [RFC9089] Xu, X., Kini, S., Psenak, P., Filsfils, C., Litkowski, S., and M. Bocci, "Signaling Entropy Label Capability and Entropy Readable Label Depth Using OSPF", RFC 9089, DOI 10.17487/RFC9089, August 2021, <<https://www.rfc-editor.org/info/rfc9089>>.
- [RFC9522] Farrel, A., Ed., "Overview and Principles of Internet Traffic Engineering", RFC 9522, DOI 10.17487/RFC9522, January 2024, <<https://www.rfc-editor.org/info/rfc9522>>.
- [RFC9790] Kompella, K., Bryant, S., Bocci, M., Mirsky, G., Ed., Andersson, L., and J. Dong, "IANA Registry and Processing Recommendations for the First Nibble Following a Label Stack", RFC 9790, DOI 10.17487/RFC9790, May 2025, <<https://www.rfc-editor.org/info/rfc9790>>.
-

Acknowledgements

This document is the result of work started in MPLS Open Design Team, with participation by the MPLS, PALS, and DETNET Working Groups.

The authors would like to thank Adrian Farrel for his contributions. The authors would also like to thank John Drake, Toerless Eckert, and Jie Dong for their comments.

Authors' Addresses

Loa Andersson

Huawei Technologies

Email: loa@pi.nu

Stewart Bryant

University of Surrey 5GIC

Email: sb@stewartbryant.com

Matthew Bocci

Nokia

Email: matthew.bocci@nokia.com

Tony Li

Juniper Networks

Email: tony.li@tony.li