

MPLS Working Group
Internet-Draft
Intended status: Informational
Expires: January 26, 2023

M. Bocci
Nokia
S. Bryant
University of Surrey 5GIC
J. Drake
Juniper Networks
July 25, 2022

Requirements for MPLS Network Action Indicators and MPLS Ancillary Data
[draft-ietf-mpls-mna-requirements-02](#)

Abstract

This document specifies requirements for MPLS network actions which affect the forwarding or other processing of MPLS packets. These requirements are derived from a number of proposals for additions to the MPLS label stack to allow forwarding or other processing decisions to be made, either by a transit or terminating LSR (i.e. the LER).

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 January 21, 2023.

Copyright Notice

Copyright (c) 2022 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Terminology	2
1.2.	Background	3
2.	Requirements Language	4
3.	MPLS Network Action Requirements	4
3.1.	General Requirements	5
3.2.	Requirements on the MNA Sub-Stack Indicator	6
3.3.	Requirements on Network Action Indicators	6
3.4.	Requirements on Ancillary Data	7
4.	IANA Considerations	8
5.	Security Considerations	9
6.	Acknowledgements	9
7.	References	9
7.1.	Normative References	9
7.2.	Informative References	9
Appendix A.	Working Group Adoption Comments	10
	Authors' Addresses	26

[1.](#) Introduction

There is significant interest in developing the MPLS data plane to address the requirements of new use cases [[I-D.saad-mpls-miad-usecases](#)], which require a general mechanism, termed MPLS network actions, for enhanced forwarding or other processing of MPLS packets. It is intended that this mechanism will be conformant to the greatest extent possible with the existing MPLS architecture as specified by, among other documents, [[RFC3031](#)], [[RFC3032](#)], and [[RFC6790](#)].

This document specifies the requirements for MPLS network actions, as well as the encoding and use of the ancillary data.

[1.1.](#) Terminology

- o Network Action: An operation to be performed on a packet. A network action may affect router state, packet forwarding, or it may affect the packet in some other way.
- o Network Action Indication (NAI): An indication in the packet that a certain network action is to be performed.

- o Ancillary Data (AD): Data in an MPLS packet associated with a given Network Action that may be used as input to the processing of the Network Action or results from the processing of the Network Action.
- o In-Stack Data: Ancillary data carried within the MPLS label stack.
- o Post-Stack Data: Ancillary data carried in a packet between the bottom of the MPLS label stack and the first octet of the user payload. This document does not prescribe whether post-stack data precedes or follows any other protocol structure such as a control word or associated channel header (ACH).
- o Scope: The set of nodes that should perform a given action.

1.2. Background

The MPLS architecture is specified in [[RFC3031](#)] and provides a mechanism for forwarding packets through a network without requiring any analysis of the packet payload's network layer header by intermediate nodes (Label Switching Routers - LSRs). Formally, inspection may only occur at network ingress (the Label edge router - LER) where the packet is assigned to a forwarding equivalence class (FEC).

MPLS uses switching based on a label pushed on the packet to achieve efficient forwarding and traffic engineering of flows associated with the FEC. While originally used for IP traffic, MPLS has been extended to support point-to-point, point-to-multipoint and multipoint-to-multipoint layer 2 and layer 3 services. An overview of the development of MPLS is provided in [[I-D.bryant-mpls-dev-primer](#)].

A number of applications have emerged which require LSRs to make forwarding or other processing decisions based on inspection of the network layer header, or some other ancillary information in the protocol stack encapsulated deeper in the packet. An early example of this was generation of a hash of the payload header to be used for load balancing over Equal Cost Multipath (ECMP) or Link Aggregation Group (LAG) next hops. This is based on an assumption that the network layer protocol is IP. MPLS was extended to avoid the need for LSRs to perform this operation if load balancing was needed based on the payload and instead use only the MPLS label stack, using the Entropy Label / Entropy Label Indicator [[RFC6790](#)] which are inserted at the LER. Other applications where the intermediate LSRs may need to inspect and process a packet on an LSP include OAM, which can make use of mechanisms such the Router Alert Label [[RFC3032](#)] or the Generic Associated Channel Label (GAL) [[RFC5586](#)] to indicate that an

intercepted packet should be processed locally. See [\[I-D.bryant-mpls-dev-primer\]](#) for detailed list of such applications.

There have been a number of new proposals for how network actions and associated ancillary data is to be carried in MPLS and how its presence is indicated to the LSR or egress LER, for example In-situ OAM and Service Function Chaining (SFC). A summary of these proposals is contained in [\[I-D.bryant-mpls-dev-primer\]](#), and an overview of use cases is provided in [\[I-D.saad-mpls-miad-usecases\]](#). [\[I-D.song-mpls-extension-header\]](#) discusses some of the issues with these proposals (note that this document draws on the requirements and issues without endorsing a specific solution from [\[I-D.song-mpls-extension-header\]](#)):

These solutions rely on either the built-in next-protocol indicator in the header or the knowledge of the format and size of the header to access the following packet data. The node is required to be able to parse the new header, which is unrealistic in an incremental deployment environment.

A piecemeal solution often assumes the new header is the only extra header and its location in the packet is fixed by default. It is impossible or difficult to support multiple new headers in one packet due to the conflicted assumption. An example of this is that the GAL/G-ACH mechanism assumes that if the GAL is present, only a single G-ACH header follows.

New applications therefore require the definition of extensions to the MPLS architecture and label stack operations that can be used across these applications in order to minimise implementation complexity and promote interoperability and extensibility.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14 \[RFC2119\]](#) [\[RFC8174\]](#) when, and only when, they appear in all capitals, as shown here.

3. MPLS Network Action Requirements

This document specifies requirements of MPLS Network Action (MNA) Indicators (NAIs), and the associated Ancillary Data, as well as the alert mechanism (the MPLS Network Action Sub-Stack Indicator) to indicate to an LSR or LER that NAIs are present in a packet. The requirements are for the behavior of the protocol mechanisms and

procedures that constitute building blocks out of which indicators for network actions and associated ancillary data are constructed. It does not specify the detailed actions and processing of any network actions or ancillary data by an LSR or LER. The purpose of this document is to identify the toolkit and any new protocol work that is required. This new protocol work **MUST** be based on the existing MPLS architecture.

3.1. General Requirements

1. MPLS combines extensibility, flexibility and efficiency by using control plane context combined with a simple data plane mechanism to allow the network to make forwarding decisions about a packet. Any solution **MUST** maintain these properties of MPLS.
2. Any MNA solutions to these requirements **MUST NOT** restrict the generality of MPLS architecture [[RFC3031](#)], [[RFC3032](#)].
3. If extensions to the MPLS data plane are required, they **MUST NOT** be inconsistent with the MPLS architecture [[RFC3031](#)], [[RFC3032](#)].
4. MNA solutions meeting the requirements set out in this document **MUST** be able to coexist with and **MUST NOT** obsolete existing MPLS mechanisms.
5. The design of any MNA solution **SHOULD** be such that an LSR is able to efficiently parse the label stack.
6. Any MNA solution specification **MUST** discuss the ECMP consequences of the design.
7. MNA solutions **MUST NOT** increase the size of the MPLS label stack more than is necessary.
8. The design of any MNA solution **MUST NOT** expose confidential information [[RFC6973](#)] [[RFC3552](#)] to the LSRs.
9. Any MNA solution specification **MUST** document any changes to the existing MPLS data plane security model that it introduces.
10. An MNA solution **MUST** allow NAI-carrying and non-NAI-carrying packets to coexist on the same LSP.

3.2. Requirements on the MNA Sub-Stack Indicator

1. An MNA solution **MUST** define how a node determines whether Network Action Indicators (NAIs) are present in the packet.
2. If NAIs are required, an MNA solution **MUST** specify where they are to be placed in the packet i.e. in-stack or post-stack.
3. An MNA solution **MUST** respect the principle that Special Purpose Labels are the mechanism of last resort and therefore must minimise the number of new SPLs that are allocated.

3.3. Requirements on Network Action Indicators

1. Insertion, parsing, processing and disposition of NAIs **SHOULD** make use of existing MPLS data plane operations.
2. An NAI **MUST NOT** be delivered to a node that is not capable of processing the indicated network action in a way that is acceptable to the imposing LER.
3. An NAI **MUST NOT** become top of stack at a node that does not understand how to perform a disposition operation on it. Disposition includes both processing and ignoring.
4. The NAI design **MUST** support scoping of network actions.
5. A given NAI specification **MUST** specify if the scope is end-to-end, hop-by-hop, or directed at one or more selected nodes.
6. If an MNA solution allows more than one scope, it must provide a mechanism **MUST** to specify the precedence of the scopes.
7. An MNA solution should support NAIs for both P2P and P2MP paths, but any specific NAI may only be supported for one or the other.
8. An MNA solution defining data plane mechanisms for NAIs **MUST** be consistent across different control plane protocol types.
9. An MNA solution **MUST** allow the in-use control / management planes to determine the ability of downstream LSRs/LERs to accept/process a given NAI.
10. An MNA solution **SHOULD** allow indicators for multiple network actions in the same packet.
11. An MNA solution **MUST** support the processing of a subset of the NAIs on a packet.

12. An MNA solution MUST enable a node inserting or modifying NAIs to determine if the far-end LER can accept and process a packet containing a given NAI.
13. NAIs are normally inserted at LERs, but MAY be processed at LSRs and LERs.
14. If a network action needs to insert an NAI with in-stack ancillary data at a transit router on an LSP, then the new network action indicator and any required ancillary data MUST be pushed onto the MPLS label stack.
15. If a network action needs to insert an NAI with below stack ancillary data at a transit router on an LSP, then the MNA solution specification MUST specify how this is achieved in all circumstances and MUST be consistent with {RFC3031}.
16. MPLS network action specifications MUST specify if in-stack or post-stack ancillary data can be rewritten by an LSR.
17. An MPLS network action specification MUST specify whether ancillary data is required and whether it is in-stack and/or post-stack.
18. Network action indicators must be allocated through the IANA process specified in the MNA solution specification.
19. An MPLS network action specification MAY support user-defined network actions. If it does, its IANA process MUST include provisions for user-defined network actions.

3.4. Requirements on Ancillary Data

1. Solutions for in-stack ancillary data MUST be able to coexist with and MUST NOT obsolete existing MPLS mechanisms. Such solutions MUST be described in a standards track RFC.
2. Specifications for MNA solutions that use in-stack ancillary data MUST justify why they require in-stack ancillary data.
3. MNA solutions MUST take care to limit the quantity of in-stack ancillary data to the minimum amount required.
4. A common preamble for ancillary data MUST be defined so that a node receiving the ancillary data can determine whether to process, ignore, skip over or discard it according to network or local policies.

5. A standardised container **MUST** be defined for in-stack ancillary data based on the MPLS LSE.
6. Any MNA solution specification **MUST** describe whether it can coexist with existing post-stack data mechanisms e.g. control words and G-ACH, and if so how this coexistence operates.
7. An MNA solution **MUST** allow an LER inserting ancillary data to determine that each node that needs to process the ancillary data can read the required distance into the packet at that node, for example [[RFC9088](#)].
8. In order to prevent unnecessary scanning of the packet, care needs to be taken in the location of post stack ancillary data, for example it **SHOULD** be located as close to the bottom of the label stack as possible.
9. Ancillary data **MAY** be associated with control or maintenance information for traffic carried by an LSP, and/or it **MAY** be associated with the user traffic itself.
10. For scoped ancillary data, any MNA solution **MUST** allow an LER inserting NAIs whose network actions make use of that ancillary data to determine if the NAI and ancillary data will be processed by LSRs within the scope along the path. Such a solution **MAY** need to determine if LSRs along the path can process a specific type of AD implied by the NAI at the depth in the stack that it will be presented to the LSR.
11. MNA solution specifications **MUST** specify if the ancillary data needs to be processed as a part of the immediate forwarding operation and whether packet mis-ordering is allowed to occur as a result of the time taken to process the ancillary data. Ed. I think this applies to both NAs and ancillary data and should be generalised.
12. A solution **MUST** be provided to verify the authenticity of ancillary data processed to LSRs [[RFC3552](#)].
13. The design of the ancillary data **MUST NOT** expose confidential information [[RFC6973](#)] [[RFC3552](#)] to the LSRs.

[4.](#) IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

5. Security Considerations

The mechanisms required by this document introduce new security considerations to MPLS. Individual solution specifications meeting these requirements MUST address any security considerations.

6. Acknowledgements

The authors gratefully acknowledge the contributions from Greg Mirsky, Yingzhen Qu, Haoyu Song, Tarek Saad, Loa Andersson, Tony Li, Adrian Farrel, Jie Dong and Bruno Decraene, and participants in the MPLS working group who have provided comments.

The authors also gratefully acknowledge the input of the members of the MPLS Open Design Team.

7. References

7.1. 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>>.
- [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>>.

7.2. Informative References

- [I-D.bryant-mpls-dev-primer]
Bryant, S., "A Primer on the Development of MPLS", [draft-bryant-mpls-dev-primer-02](#) (work in progress), May 2022.
- [I-D.saad-mpls-miad-usecases]
Saad, T., Makhijani, K., Song, H., and G. Mirsky, "Use Cases for MPLS Network Action Indicators and MPLS Ancillary Data", [draft-saad-mpls-miad-usecases-02](#) (work in progress), April 2022.
- [I-D.song-mpls-extension-header]
Song, H., Li, Z., Zhou, T., Andersson, L., Zhang, Z., Gandhi, R., Rajamanickam, J., and J. Bhattacharya, "MPLS Post-Stack Extension Header", [draft-song-mpls-extension-header-07](#) (work in progress), July 2022.

- [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>>.
- [RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", [BCP 72](#), [RFC 3552](#), DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/info/rfc3552>>.
- [RFC4221] Nadeau, T., Srinivasan, C., and A. Farrel, "Multiprotocol Label Switching (MPLS) Management Overview", [RFC 4221](#), DOI 10.17487/RFC4221, November 2005, <<https://www.rfc-editor.org/info/rfc4221>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", [RFC 5586](#), DOI 10.17487/RFC5586, June 2009, <<https://www.rfc-editor.org/info/rfc5586>>.
- [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>>.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", [RFC 6973](#), DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/info/rfc6973>>.
- [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>>.

[Appendix A](#). Working Group Adoption Comments

1. Normative Language

Use of the normative language, Terminology section in particular. For example, in "There may be associated ancillary data in the packet."

Ed.> We will check all of the normative language, but in this instance it is not needed.

2. The NAS abbreviation

Network Action Sub-Stack is abbreviated as NAS. I think that abbreviating as NASS or presenting the extended term as Network Action Sub-stack improves correlation between the full form and acronym.

Ed.> Agreed. Will change to NASS throughout.

3. non-use of NAS abbreviation

Also, I've noticed that NAS is not used throughout the document. It might not be useful after all.

Ed.> Will check throughout and remove if not used.

4. [Section 3.2](#) typo

Perhaps a typo in the first requirement [Section 3.2](#) s/and/an/ It is not clear what "NAI is delivered to a node" might mean in the requirement 5 of [Section 3.2](#). Perhaps the next requirement is sufficient and #5 can be removed from the document.

Ed.> OK. Will check

5. Extra words

Also, #5 seems like it has some extra wording. Perhaps s/in the way in a way/in a way/?

Ed.> OK will check

6. Merging of drafts?

One thing I'm debating is whether [draft-bryant-mpls-dev-primer-01](#) - A Primer on the Development of MPLS (ietf.org) should be merged with this draft?

Ed.> Disagree. The primer draft is a useful background source, but it is not a set of requirements.

7. ADI -> NAI

In this version the term "Ancillary Data Indicator" is changed to "Network Action Indicator". While there is some difference between the definition of the two terms: Ancillary Data Indicator (ADI): A indicator in the MPLS label stack that ancillary data exists in this packet. It MAY also indicate the specific type of the ancillary data. Network Action Indication (NAI): An indication in the packet that a certain network action is to be performed. There may be associated ancillary data in the packet. The above definition shows that ADI firstly is the indicator of the existence of the ancillary data, and optionally can be the indicator of specific type of ancillary data. While NAI is only the indicator of a certain type of network action. Thus NAI cannot replace ADI directly in this document. I'd suggest to add the ADI back to the terminology section, and change all the NAI in [section 3.2](#) back to ADI. If needed, the requirements on NAI can be added as separate items.

Ed.> Resolved

8. Addition to [section 3.1](#)

For backward compatibility and consistency, It is suggested to add the below items to [section 3.1](#) as general requirements: a) Solutions meeting the requirements set out in this document MUST be compatible with existing MPLS mechanisms. b) Solutions meeting the requirements set out in this document MUST reuse existing MPLS mechanisms when possible. c) For network actions which are developed or under development in IETF, the encoding and processing of the network action data MUST be reused.

Ed.> Requirements 1-4 already cover points a and b. C is outside the scope of this document but we are not intending to limit or prescribe anything about existing standards.

9. Action Indicators without AD

Do not use the ADI for Network Actions that does not have ancillary data, use NADI (non-ADI).

Ed.> Discuss with John

10. AD definition

I was under the impression reading Jie's note that actions itself are the Ancillary Data. Your definition of "Ancillary

Data" seems to be limited to action parameters or metadata which is likely why you draw such conclusions.

Ed.> Agree to clarify AD is associated with an NA (see John's text)

11. Optional AD?

AD definition treats anything new to the current label stack as an optional add-on == ancillary while NAI treats only optional parameters or metadata associated with newly defined actions as ancillary.

Ed.> See proposed text.

12. ADI and NAI are different

It is also my understanding that the definition of ADI and NAI are different. ADI is used to indicate that there is information in addition to the legacy label stack in the packet, while NAI is used to indicate a certain type of network action. The existence of NAI and the optional data associated with the action need an indicator, which is the ADI.

Ed.> See proposed text.

13. Retire ADI

There is no need for an ancillary data indicator and we should probably retire the term.

Ed.> Agree ADI will be removed.

14. AD generic?

As Robert and I mentioned, the term ancillary data is generic and refers to all types network actions information, including those with and without additional action data. Thus NAI can be considered as one type of ancillary data.

Ed.> Ancillary data may be explicit in the packet or inferred from context. See the new definition in the terminology section.

15. What ADI indicates

An ADI is the indicator of the presence of ANY non-label information in the MPLS packet. Following it there may be

indicators of each specific network action. And my understanding is the requirements in [section 3.2](#) was mainly on the ADI.

Ed.> ADI will be removed

16. Common requirement to carry AD

According to the use cases collected, their common requirement is to carry ancillary data in the data packet to affect the packet forwarding or processing behavior, or the network states. There is no specific requirement on where the ancillary data should be put in the packet. Thus in the requirement document it would be clear enough to just mention ancillary data and its indicator (ADI).

Ed.> Addressed with a requirement that a standardised container for in-stack data is required.

17. Not discussed

We have not discussed changing ADI to NAI.

18. Discussed when draft presented in the Open DT

The change of ADI to NAI was presented when the new version of the Framework were discussed in the Open DT.

19. Comment from Loa

I think both comment 17 and 18, could be misunderstood, yes the NAI was introduced (for good reason) in new Framework, however it does not in itself exclude have an ADI in a packet, only that is not the all-emcompassing term that it was earlier.

Ed.> Needs new text that explains that NAI implies that there may be ancillary data to process.

20. Keep using ADI

My suggestions would be to keep using ADI for the generic indicator, and could use NAI for the indicator of specific action. I don't think we need to mention NAS here, which is specific to ISD based solution.

Ed.> Ancillary data should only be included if there is some action associated with it. Presence of AD is implicit in the

NAI definition. A standardised container for AD could allow empty data.

21. Use of NSI

I propose Network Action Sub-stack Indicator (NSI) for this purpose.

Proposed definition: An LSE used to indicate the presence of a Network Action Sub-stack.

Ed.> NAS and NASI removed

22. Revise definition of NAS

We should also revise the definition of NAS to use this (21).

Ed.> See framework

23. Popping NAS

When you pop a stack in programming, the concept that MPLS borrowed, you pop the procedure indicator and the procedure parameters. This is consistent with popping an NAS

Ed.> No action

24. More than a name change

No it is more than just a name change. There is a concept change and we re-wrote a bunch of text to align with the NAI approach. For example an NAI may not have AD to indicate.

Ed.> Yes

25. Writeable NAS

I have some further questions. NAS is something within the label stack but writable by intermediate nodes. Is this the stack operation? Besides, if NAS emerges at ToS, you said it'll be popped and discarded. What if the NAS also needs to be applied to the labels below it? Whatever measures you will take here, are those the stack operations?

Ed.> NAS has been removed.

26. The NFRR use case

I have several questions about the NFRR use case. As I understand it, a point of local repair (PLR) imposes NAS with the NFRR indicator so that it becomes ToS at the merge node (MN). If that is correct, then the MN will remove the NAS with the NFRR indicator as the packet is returned on the "normal" path. Hence, I don't see why an intermediate node would need to write into an existing in the label stack NAS in support of NFRR.

Ed.> NFRR is not discussed in this draft.

27. Intermediate node re-writing

I may not see a case for an intermediate node re-writing the existing NAI. I think that the node should impose a new NAS. I see the case for an intermediate node writing into PSD (e.g., HbH IOAM though I think that is too expensive and the IOAM Direct Export or Hybrid Two-Step are a better choice). To summarize, I don't see the need for an intermediate node to write into ISD and am open to discussing the node writing into PSD.

Ed.> Added a requirement that network action specifications must describe whether ISD can be rewritten by an intermediate node.

28. Two types of indicators

One of my concern is that there are two types of indicators, and they cannot be represented using the same term. - The first one (call it indicator-1) is used to indicate the presence of any non-forwarding-label information in the packet. As discussed it may be realized as a bSPL. It does not indicate the type of actions to be performed. - The second type of indicator (call it indicator-2) is used to indicate the presence of a specific type of action. Such action may or may not have associated data with it. This may be realized as a Flag or an action type in the packet.

Ed> Not sure this is correct. The bSPL indicates the presence of NAIs. The NAI specifies the network action and whether ancillary data is present, and where it is present. This should be clarified in the framework.

29. Second term needed

do not understand why we need another term. We have not had a situation where we wanted to refer to "NAS + PSD" and lacked a term for it.

Ed.> Agreed

30. Propose text

That said, please feel free to propose a term and a definition. I do not understand why we need another term. We have not had a situation where we wanted to refer to "NAS + PSD" and lacked a term for it.

That said, please feel free to propose a term and a definition.

Ed.> Agreed

31. NAS not general enough

Thus NAS is not general enough to cover the PSD. What we need is a generic term to cover all the possible cases of ancillary data. Furthermore, the indicator and the ancillary data would need separate terms anyway and does not need to be coupled under one term.

Ed.> NAS has been removed. The definition of Network Action should cover this.

32. Use of ancillary data

- * We are already using ancillary data for this purpose
- * Exactly, and that's why we don't need to mention NAS in the requirement document.

Ed> NAS has been removed.

33. What the NAS contain

- * Let me put it this way: By definition, NAS contains: ADI, Optional NAI and Optional ISD. While what we want to refer to with the term is: Optional NAI Optional ISD and Optional PSD. Note it does not include the ADI.
- * This is incorrect. It contains a network action sub-stack indicator, network action indicators, and any in-stack ancillary data (as defined by the specified network actions)

Ed.> NAS has been removed.

34. User-defined actions

user-defined network actions? Should we mention in the requirements doc?

Ed.> The allocation policy should be specified in the draft that defines the IANA registry for the NAIs.

35. [draft-ietf-mpls-mna-requirements](#)

I hope, if adopted, the filename can be adjusted to use MRN not MIAD-ADI.

Comment from Loa: I think the filename we are considering is: [draft-ietf-mpls-mna-requirements](#)

Ed.> Done.

36. In the Abstract (1)

This work is the product of the IETF MPLS Open Design Team.

Before posting as a Working Group draft, this sentence needs to be removed. It's OK to say something in the Acknowledgements.

Loa: Depending on what is meant by "before", I have a comment on this, just because info is at the wrong place it should not be considered "blocking" and could be update at any time. Personally I think working group draft version -01 would be a good place to do this update.

Ed> OK agreed.

37. In the Abstract (2)

The term "application data" may be (is) confusing. While you probably mean it to imply an application of MPLS, it may be confused with the type of application that runs end-to-end (i.e., on a host). Although, reading some of the text, it does feel like some of the time you are intending to imply that the application software may somehow be able to provide the ancillary data that is ultimately carried by MPLS.

I think, in the Abstract, you could say...

based on ancillary data that may be carried in or below the bottom of the label stack.

...but you should probably look through the rest of the text and consider the use of the word "application." Maybe one approach

would be to specifically call out the term near the top of the document in order to correctly set the context.

Ed> OK agreed.

38. In [section 1](#) (1).

is then processed using mechanisms implemented by intermediate and/or egress LSRs that comply with the MPLS base architecture and potentially its extensions, including (but not limited to) [[RFC3031](#)], [[RFC3032](#)], [[RFC6790](#)].

This sounds like the mechanisms to process the ancillary data are already specified in those RFCs, but (of course) that's not the case.

You are probably making a point about the nodes being "MPLS" and also saying something about backward compatibility of the mechanism. But it should be clear that only nodes that contain new functionality will be able to recognise and process ancillary data.

Ed.> OK will fix

39. In [section 1](#) (2).

This draft specifies

Say 'document' so that you are future-proofed for publication as an RFC.

Ed.> OK will fix.

40. In [section 1.1](#) (1)

s/an Label/a Label/ s/perfomed/performed/

Ed.> OK will fix

41. In [section 1.1](#) (2).

* Network Action: An operation to be performed on a packet. A network action may affect router state, packet forwarding, or it may affect the packet in some other way.

If the operation affects router state, is it really performed on the packet?

Ed.> Yes.

42. In [section 1.1](#) (3)

- * Network Action Indication (NAI): An indication in the packet that a certain network action is to be performed. There may be associated ancillary data in the packet
- * Network Action Sub-Stack (NAS): A set of related, contiguous Label Stack Entries (LSEs). The first LSE contains the NAI. The TC and TTL values in the sub-stack may be redefined.

The first bullet simply says that the NAI is "in the packet", but the second bullet goes on to define where/how it is carried. I would say that it is totally irrelevant to the `_requirements_` how the NAI is encoded/carried (although there may be some requirements that limit the options). But I note that there is no further mention of the NAS or of a "sub-stack". I suggest removing this second bullet.

Ed.> Agreed. Will fix.

43. In [section 1.1](#) (4)

- * In-Stack Data: Any data within the MPLS label stack including the outer LSE and the bottom of stack (the LSE with the S-bit set).
- * Post-Stack Data: Any data beyond the LSE with the S-Bit set, but before the first octet of the user payload. This document does not prescribe whether post-stack data precedes or follows any other protocol structure such as a control word or associated channel header (ACH).

Does "any data" mean "any ancillary data"?

Ed.> Agreed. Will fix.

44. In [section 1.2](#) (1)

s/number of new proposals/number of proposals/

Ed.> Agreed

45. In [Section 1.2](#) (2)

for example In-situ OAM and Service Function Chaining (SFC)

Might benefit from references for iOAM and SFC.

Ed.> Will add references

46. In [section 1.2](#) (3)

[I-D.song-mpls-extension-header] summarises some of the issues with existing solutions to address these new applications (note that this document draws on the requirements and issues without endorsing a specific solution from [\[I-D.song-mpls-extension-header\]](#)):

This gives more emphasis to the referenced draft than I think you intend. If you intend that people read that draft to see the issues, it is a normative reference. But if you are just mentioning it and have pulled the information into this document, then you need to reduce the emphasis. How about...

[I-D.song-mpls-extension-header] sets out some of the issues in how existing solutions address these new applications. This document draws on the requirements and issues noted in that document without endorsing any specific solution.

Ed.> New text is virtually identical. The existing text is fine and is not intended to be normative.

47. [Section 2](#)

While I understand the desire to express the requirements in definitive language, [BCP14](#) is not about requirements. Rather it is intended to describe implementation behaviours.

A way around this that is often used is to include a subsequent paragraph such as...

Although this document is not a protocol specification, this convention is adopted for clarity of description of requirements.

See, for example, [RFC 4139](#), [RFC 4687](#), or [RFC 5862](#).

Ed> Despite what [BCP14](#) says, there is a long standing convention of using must and should in requirements documents because they specify requirements on protocol design. Suggest no change to the existing text.

48. In [section 3.2](#) (1)

s/and indicator/an indicator/

Ed> OK

49. In [section 3.2](#) (2)

(2). An MPLS Network Action MUST specify whether ancillary data is required in the label stack and/or post-stack data.

Do you mean this, or do you mean that this must be specified in the documentation of the NA?

Ed.> OK will fix

50. In [section 3.2](#) (1)

(3). Any solution MUST respect the principle that Special Purpose Labels are the mechanism of last resort and therefore must minimise the number of new SPLs that are allocated.

Presumably a minimum here would be zero?

Loa: No we have considered this and decided that there is room to specify 1 bSPL for MNA.

Ed.> Existing text is fine.

Loa: I also think that we should s/Special Purpose Labels/Base Special-Purpose Labels Note: for the Extended SPLs there are no such restriction.

Ed.> Existing text is fine.

51. In [section 3.2](#) bullet 5

s/in the way in a way/in a way/

Ed.> OK will fix

52. In [section 3.2](#) (bullet, 5, 6 and 10)

Bullet 10 is a wholly contained subset of bullet 5. Actually, bullet 10 is a wholly contained subset of bullet 6. Makes me think that bullets 5 and 6 possibly say the same thing as each other.

Ed.> Will move 10 up the list to be closer to 5 / 6. They are not the same but are related.

53. In [section 3.2](#) (2)

(11). NAIs SHOULD be supported for both P2P and P2MP paths, but any specific NAI may only be supported for one or the other.

Really? You can't have an NAI that is equally applicable for both P2P and P2MP? Seems an odd restriction to impose.

Ed.> The text is correct as written.

54. In [section 3.2](#) (3)

(15). NAIs can only be inserted at LERs, but MAY be processed at LSRs and LERs. If it is required to insert an NAI at a transit LSR on an LSP, then a new label stack MUST be pushed.

What does it mean to push a new label stack? If you mean that we should support "MPLS in MPLS" encapsulation so that the packet has two bottom of stack bits set, I should point out that this was previously discussed and abandoned because the presence of an LSE immediately after a set bottom of stack bit was considered unacceptable because of the assumptions made by existing hardware about what follows the bottom of stack.

55. In [section 3.2](#) (4)

(19). Any specification of a solution that inserts or modifies the NAI MUST discuss the possible ECMP consequences.

This seems to at least partially contradict 3.2/15

It is also not clear what it means "to modify an indication in the packet that a certain network action is to be performed". I guess it means to remove the NAI?

Ed.> See updated text

56. In [section 3.3](#) (1)

(3.3/1) is surely already covered by 3.3/4

Ed>These are correct but we have strengthened 3.3/1.

57. In [section 3.3](#) (2)

(3.3/3) seems to be unnecessary given 3.3/1

Ed.> (1) is in-stack, (3) is post-stack, so they are correct as written.

58. Semantic Routing

I think the proposal here falls in the scope of "Semantic Routing". That is, adding information to packets so that the forwarding decisions may be enhanced to act not just on the destination address or next hop label, but also on the additional information. The precise forwarding action may be known by the forwarders by definition (such as a protocol specification), installed by a routing engine according to a routing algorithm acting on information exchanged by routing protocols, or programmed into the forwarder from a management or orchestration system.

We wrote an introduction to the idea of Semantic Routing [draft-farrel-irtf-introduction-to-semantic](#) -routing which you can look at if you want some context.

We also set out to examine the challenges and concerns introduced by Semantic Routing in [draft-king-irtf-challenges-in-routing](#) and I think it would be good if this work was calibrated against those challenges.

Loa: While semantic routing is interesting, and good to be "calibrated" against MNA as a whole (guiding documents and solutions), I think it is out of scope for the requirement document.

Ed.> Agree with Loa's response.

59. Understanding of Use Cases

I would think that a more detailed understanding of the use cases is needed before moving ahead with the requirements. I wouldn't go as far as saying that the use cases need to be referenced normatively, but I do think they need a little more attention from the WG to motivate actually adopting this work. That is, this document shows what we might need to do, but without the use cases, we would be doing it "because we can" and "because it might be useful one day." Those are not, I think really good reasons to make fairly substantial changes to deployed forwarding paradigms.

This is not to say that I dispute that there may be some valuable use cases, but that the WG needs to agree which ones

are important in order to be sure that the requirements are on target.

Ed.> The use cases identify what is currently being discussed and are driving this activity. The requirements reflect the current state of the use cases. The use cases are necessary but not sufficient to fully derive the requirements.

60. Conflicting text in document

I'm puzzled that some of the text in this document appears to limit itself to cases that require ancillary data, while other parts also consider the requirements for network functions that don't require ancillary data, but do still need to be encoded in the label stack in some way. I suspect this is just editorial, but while the document title is "Requirements for MPLS Network Action Indicators and MPLS Ancillary Data" the Abstract says "This draft specifies requirements for indicators in the MPLS label stack to support ancillary data in the packet and high level requirements on that ancillary data," and the Introduction seems entirely focused on the ancillary data case.

It would be good to be clear, at the point of adoption, which way we are jumping on this question.

[draft-andersson-mpls-mna-fwk](#)

seems to be fully behind network actions some of which may also require ancillary data.

Perhaps this document should reference that one for additional information?

Ed> Updated abstract to focus on network actions. Will check for inconsistencies in rest of document.

61. Loa: I have been thinking about a short text (1 or 2 paragraphs) on how the guiding documents fit together that should appear, e.g. in the introduction of all 3 documents. It could be added later in the process, but should be there when the documents go to wglc. Let me know if there are someone that will help work on this,

Ed.> Agreed. Will add to todo list.

Authors' Addresses

Matthew Bocci
Nokia

Email: matthew.bocci@nokia.com

Stewart Bryant
University of Surrey 5GIC

Email: sb@stewartbryant.com

John Drake
Juniper Networks

Email: jdrake@juniper.com

