MPLS Working Group
Internet-Draft
Intended status: Informational
Expires: 17 February 2024

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

16 August 2023

Requirements for MPLS Network Actions
draft-ietf-mpls-mna-requirements-06

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 17 February 2024.

Table of Contents

1. Introduction 2
1.1. Terminology 2
1.2. Background 3
2. Requirements Language 5
3. MPLS Network Action Requirements 5
3.1. General Requirements 5
3.2. Requirements on the MPLS Network Action Sub-Stack Indicator 6
3.3. Requirements on Network Action Indicators 6
3.4. Requirements on Ancillary Data 8
4. IANA Considerations 9
5. Security Considerations 9
6. Acknowledgements 9
7. References 9
7.1. Normative References 9
7.2. Informative References 10
Authors' Addresses 12

1. Introduction

There is significant interest in developing the MPLS data plane to address the requirements of new use cases [I-D.ietf-mpls-mna-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

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

* Network Action Indication (NAI): An indication in the packet that a certain network action is to be performed.

* 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.

* In-Stack Data: Ancillary data carried within the MPLS label stack.

* 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).

* 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 [I-D.gandhi-mpls-ioam-sr] and Service Function Chaining (SFC) [RFC7665]. 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.ietf-mpls-mna-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 use cases therefore require the definition of extensions to the MPLS architecture and label stack operations that can be used across these use cases 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.

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

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] and [RFC5331].

3. If extensions to the MPLS data plane are required, they MUST NOT be inconsistent with the MPLS architecture [RFC3031], [RFC3032] and [RFC5331].

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. MNA solutions that increase the size of the MPLS label stack in a way that is not controlled by the ingress LER MUST discuss the consequences.

9. The design of any MNA solution MUST NOT expose confidential information [RFC6973] [RFC3552] to the LSRs.

10. Any MNA solution specification MUST document any changes to the existing MPLS data plane security model that it introduces.

11. An MNA solution MUST allow NAI-carrying and non-NAI-carrying packets to coexist on the same LSP.

3.2. Requirements on the MPLS Network Action 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 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.

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 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. NAIs SHOULD only be inserted at LERs, and MAY be processed at LSRs and LERs.

13. If a network action needs to insert an NAI with in-stack ancillary data at an LSR on an LSP, then the new network action indicator and any required ancillary data MUST be pushed onto the MPLS label stack.

14. If a network action needs to insert an NAI with below stack ancillary data at an LSR on an LSP, then the MNA solution specification MUST specify how this is achieved in all circumstances and MUST be consistent with {RFC3031}.

15. MPLS network action specifications MUST specify if in-stack or post-stack ancillary data can be rewritten by an LSR.

16. An MPLS network action specification MUST specify whether ancillary data is required and whether it is in-stack and/or post-stack.

17. Network action indicators MUST be allocated through the IANA process specified in the MNA solution specification.

18. Is it RECOMENDED that an MPLS network action specification supports network actions for private use [RFC8126].

19. A node removing an NAI MUST NOT leave the MPLS label stack in such a way that downstream nodes are unable to determine the presence of ancillary data remaining in the packet.

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 any 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. We 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.

14. A mechanism MUST exist to notify an egress LER of the presence of ancillary data so that it can dispose of it appropriately.

15. An egress LER MUST NOT forward a packet with ancillary data to a node that is not expecting the ancillary data to be present.

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/rfc/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

7.2. Informative References

[I-D.bryant-mpls-dev-primer]
Bryant, S., "A Primer on the Development of MPLS", Work in Progress, Internet-Draft, draft-bryant-mpls-dev-primer-02, 9 May 2022, <https://datatracker.ietf.org/doc/html/draft-bryant-mpls-dev-primer-02>.

[I-D.gandhi-mpls-ioam-sr]
Gandhi, R., Ali, Z., Filsfils, C., Brockners, F., Wen, B., and V. Kozak, "MPLS Data Plane Encapsulation for In-situ OAM Data", Work in Progress, Internet-Draft, draft-gandhi-mpls-ioam-sr-06, 18 February 2021, <https://datatracker.ietf.org/doc/html/draft-gandhi-mpls-ioam-sr-06>.

[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-02, 13 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-usecases-02>.

[I-D.song-mpls-extension-header]
Song, H., Zhou, T., Andersson, L., Zhang, Z. J., and R. Gandhi, "MPLS Network Actions using Post-Stack Extension Headers", Work in Progress, Internet-Draft, draft-song-mpls-extension-header-12, 14 April 2023, <https://datatracker.ietf.org/doc/html/draft-song-mpls-extension-header-12>.

[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/rfc/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/rfc/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/rfc/rfc3552>.

[RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, DOI 10.17487/RFC5331, August 2008, <https://www.rfc-editor.org/rfc/rfc5331>.

[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/rfc/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/rfc/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/rfc/rfc6973>.

[RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", RFC 7665, DOI 10.17487/RFC7665, October 2015, <https://www.rfc-editor.org/rfc/rfc7665>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/rfc/rfc8126>.

[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/rfc/rfc9088>.

Authors' Addresses

Matthew Bocci (editor)
Nokia
Email: matthew.bocci@nokia.com

Stewart Bryant
University of Surrey 5GIC
Email: sb@stewartbryant.com

John Drake
Juniper Networks
Email: jdrake@juniper.com