

MPLS Working Group
Internet-Draft
Intended status: Informational
Expires: April 17, 2022

M. Bocci
Nokia
S. Bryant
University of Surrey 5GIC
October 14, 2021

Requirements for MPLS Label Stack Indicators for Ancillary Data
draft-bocci-mpls-miad-adi-requirements-00

Abstract

This draft specifies requirements for indicators in the MPLS label stack of ancillary data that exists below the label stack. This work is the product of the IETF MPLS Open Design Team. Requirements are derived from a number of new 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, based on application data that may be in or below the bottom of the label stack.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 April 17, 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

MIAD ADI Requirements

October 2021

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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.	MPLS Ancillary Data Indicator Requirements	4
2.1.	General Requirements	4
3.	IANA Considerations	7
4.	Security Considerations	7
5.	Acknowledgements	7
6.	References	7
6.1.	Normative References	7
6.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

There is significant interest in developing the MPLS data plane to address the requirements of new applications. These applications typically include ancillary data that is contained in or below the label stack. There is a requirement for this data to be either intercepted and processed, or some other forwarding decision to be made. This makes use of mechanisms implemented by an intermediate or egress label switching router (LSR) that complies with the MPLS base architecture and potentially its extensions, including (but not limited to) [[RFC3031](#)], [[RFC3032](#)], [[RFC6790](#)].

This draft specifies requirements for indicators in the MPLS label stack to support these applications.

[1.1.](#) Terminology

- o Ancillary Data: Data relating to the MPLS packet that may be used

to affect the forwarding or other processing of that packet, either at the LER or LSR. This data may be implicit (i.e. context-specific), encoded within the label stack (in-stack data), after the bottom of the label stack but not considered a part of the payload, or within the payload.

- o In-Stack: Any location within the MPLS label stack including the outer label and the bottom of stack (the label with the S-bit set).
- o 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.

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 ancillary data is 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](#)], an overview of use cases is provided in [Reference to MIAD use cases]. [[I-D.song-mpls-extension-header](#)] summarises some of the issues with existing solutions to address these new applications:

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.

[2.](#) MPLS Ancillary Data Indicator Requirements

This document specifies requirements of MPLS Indicators for Ancillary Data (MIAD). The requirements are for the behavior of the protocol mechanisms and procedures that constitute building blocks out of which mechanisms for indicating ancillary data that exists in

the MPLS payload using the MPLS label stack (so-called in-stack indicators) are constructed. It does not specify the detailed processing that may be required by an application of that ancillary data by an LSR. The requirements in this document do not describe what functions MIAD implementation supports. 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.

[2.1.](#) General Requirements

- o 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.
- o Any solutions to these requirements MUST not restrict the generality of MPLS architecture.

- o Any solution MUST respect the principle that Special Purpose Labels are the mechanism of last resort.
- o Solutions MUST be able to coexist with and not obsolete existing MPLS mechanisms.
- o Neither an ADI or ancillary data must be delivered to a node that is not capable of processing it.
- o Care needs to be taken in the coexistence of ancillary data and existing post-stack data mechanisms.
- o Mechanisms are required to determine that all nodes that need to process the ancillary data can read the required distance into the packet at that node.
- o A mechanism is REQUIRED for Ancillary Data Indicators to indicate the presence of ancillary data in the MPLS label stack (Ed. note: This is similar to ELI).
- o A mechanism is REQUIRED for Ancillary Data Indicators to indicate the presence of ancillary data below the MPLS label stack (Ed.

note: this is similar to GAL/G-ACH).

- o The mechanism to indicate that Ancilliary Data is present MUST operate in the context of the top of stack LSE.
- o Ancilliary data may be associated with control or maintenance information for traffic carried by an LSP, or it may be associated with the user traffic itself.
- o Ancilliary Data Indicators (ADIs) SHOULD make use of existing MPLS data plane operations. If extensions to the MPLS data plane are required, they MUST NOT be inconsitent with the MPLS architecture.
- o A mechanism is REQUIRED to enable an LER inserting ADIs to determine whether LSRs along the path can parse the label stack and process the ADI at the location it is inserted.
- o A mechanism is REQUIRED to enable an LER inserting ADIs to determine if the ADI will be processed by LSRs along the path.
- o A mechanism is REQUIRED to enable an LER inserting ADIs to determine if the far-end LER can accept and process a packet containing a given ADI.
- o ADIs SHOULD be supported for both P2P and P2MP paths, but any specific ADI may only be supported for one or the other.

- o Data plane mechanisms for ADIs MUST be independent of the control plane type (LDP, RSVP, BGP, static, IGP, etc).
- o A mechanism MUST be defined for control planes (LDP, RSVP, BGP, static, IGP, etc) to determine the ability of downstream LSRs/LEs to accept/process a given ADI.
- o A mechanism is REQUIRED to enable an LSR to efficiently determine if an ADI is present in a packet.
- o ADIs can only be inserted at LERs, but may be processed at LSRs and LERs. If it is required to insert an ADI at a transit router on an LSP, then a new label stack must be pushed. .
- o It SHOULD be possible to include indicators for ancillary data for

multiple applications in the same packet, but each ADI only supports one application.

- o It MUST be possible to insert new ADIs for new applications on the same LSP.[Ed note: need to clarify]
- o The solution must allow ADI and non-ADI packets to coexist on the same LSP.
- o The solution must support the processing of a subset of the ADIs on a packet.
- o The solution MUST support slow path processing of ancillary data.
- o The solution MUST support fast path processing of ancillary data.
- o The solution MUST support hop-by-hop processing of ancillary data.
- o The solution MUST support end-to-end processing of ancillary data.
- o If both hop-by-hop and end-to-end ancillary data indicators are present together, the precedence must be specified in the design.
- o In order to prevent unnecessary scanning of the packet, care needs to be taken in the location of the ancillary data, for example it should be located as close to the label stack as possible.
- o A solution must be provided to verify the authenticity of ancillary data processed to LSRs.

- o The design of the ADIs and ancillary data must not expose confidential information to the LSRs.

[3.](#) IANA Considerations

This document makes no request of IANA.

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

RFC.

[4.](#) Security Considerations

The mechanisms required by this document introduce new security considerations to MPLS. It is expected that individual solutions meeting these requirements will address any security considerations.

[5.](#) Acknowledgements

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

[6.](#) References

[6.1.](#) 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>>.

[6.2.](#) Informative References

[I-D.bryant-mpls-dev-primer]
Bryant, S., "A Primer on the Development of MPLS", [draft-bryant-mpls-dev-primer-00](#) (work in progress), March 2021.

[I-D.song-mpls-extension-header]
Song, H., Li, Z., Zhou, T., Andersson, L., and Z. Zhang, "MPLS Extension Header", [draft-song-mpls-extension-header-05](#) (work in progress), July 2021.

[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>>.

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>>.

[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>>.

Authors' Addresses

Matthew Bocci
Nokia

Email: matthew.bocci@nokia.com

Stewart Bryant
University of Surrey 5GIC

Email: sb@stewartbryant.com