Workgroup: MPLS Working Group
Internet-Draft:
draft-saad-mpls-miad-usecases-00
Published: 7 January 2022
Intended Status: Informational
Expires: 11 July 2022
Authors: T. Saad K. Makhijani
Juniper Networks Futurewei Technologies
H. Song
Futurewei Technologies

Usecases for MPLS Indicators and Ancillary Data

Abstract

This document presents a number of use cases that have a common need for encoding MPLS function indicators and ancillary data inside MPLS packets. The use cases described are not an exhaustive set, but rather the ones that are actively discussed at the MPLS Working Group.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

This document describes important cases that require carrying additional ancillary data within the MPLS packets, as well as the means to indicate ancillary data is present.

These use cases have been identified by the MPLS working group design team working on defining MPLS function indicators and ancillary data for the MPLS data plane. The use cases described in this document will be used to assist in identifying requirements and issues to be considered for future resolution by the working group.

* ID: draft-gandhi-mpls-ioam describes applicability of IOAM to MPLS dataplane.

* RFC 8986 describes the network programming usecase for SRv6 dataplane.

*RFC 8595 describes solution for MPLS-based forwarding for Service Function Chaining

1.1. Terminology

The following terminology is used in the document:

IETF Network Slice:
a well-defined composite of a set of endpoints, the connectivity requirements between subsets of these endpoints, and associated requirements; the term 'network slice' in this document refers to 'IETF network slice' as defined in [I-D.ietf-teas-ietf-network-slices].

IETF Network Slice Controller (NSC):
controller that is used to realize an IETF network slice [I-D.ietf-teas-ietf-network-slices].

Network Resource Partition:
the collection of resources that are used to support a slice aggregate.

Time Sensitive Networking:
Networks that transport time sensitive traffic.

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.

1.2. Acronyms and Abbreviations

MIAD: MPLS Label Stack Indicators for Ancillary Data

ISD: In-stack data

PSD: Post-stack data

MPLS: Multiprotocol Label Switching

2. Use Cases

2.1. In-situ OAM

In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information within the packet while the packet traverses a particular path in a network domain.

The term "in-situ" refers to the fact that the IOAM data fields are added to the data packets rather than being sent within the probe packets specifically dedicated to OAM or Performance Measurement (PM).

IOAM can run in two modes End-to-End (E2E) and Hop-by-Hop (HbH). In E2E mode, only the encapsulating and decapsulating nodes will process IOAM data fields. In HbH mode, the encapsulating and decapsulating nodes as well as intermediate nodes process IOAM data fields.

The IOAM data fields are defined in [I-D.ietf-ippm-ioam-data], and can be used for various use-cases of OAM and PM.

[I-D.gandhi-mpls-ioam-sr] defines how IOAM data fields are transported using the MPLS data plane encapsulations, including Segment Routing (SR) with MPLS data plane (SR-MPLS).

IOAM data are added after the bottom of the label stack. The IOAM data fields can be of fixed or incremental size as defined in [I-D.ietf-ippm-ioam-data]. [I-D.gandhi-mpls-ioam] describes applicability of IOAM to MPLS dataplane. The encapsulating MPLS node needs to know if the decapsulating MPLS node can process the IOAM data before adding it in the packet.

2.2. Network Slicing

[I-D.ietf-teas-ietf-network-slices] specifies the definition of a network slice for use within the IETF and discusses the general framework for requesting and operating IETF Network Slices, their characteristics, and the necessary system components and interfaces.

Multiple network slices can be realized on top of a single shared network.

In order to overcome scale challenges, IETF Network Slices may be aggregated into groups according to similar characteristics. The slice aggregate [I-D.bestbar-teas-ns-packet] is a construct that comprises of the traffic flows of one or more IETF Network Slices of similar characteristics.

A router that requires forwarding of a packet that belongs to a slice aggregate may have to decide on the forwarding action to take based on selected next-hop(s), and the forwarding treatment (e.g., scheduling and drop policy) to enforce based on the associated per-hop behavior.

The routers in the network that forward traffic over links that are shared by multiple slice aggregates need to identify the slice aggregate packets in order to enforce the associated forwarding action and treatment.

An IETF network slice need MAY support the following key features:

1. A Slice Selector

2. A Network Resource Partition associated with a slice aggregate.

3. A Path selection criteria

4. Verification that per slice SLOs are being met. This may be done by active measurements (inferred) or by using IOAM.

5. Additionally, there is an on-going discussion on using Service Functions (SFs) with network slices. This may require insertion of an NSH.

6. For multi-domain scenarios, a packet that traverses multiple domains may encode different identifiers within each domain.

2.2.1. Global Identifier as Slice Selector

A Global Identifier as a Slice Selector (GISS) can be encoded in the MPLS packet as defined in [I-D.kompella-mpls-mspl4fa], [I-D.li-mpls-enhanced-vpn-vtn-id], and [I-D.decraene-mpls-slid-encoded-entropy-label-id]. The Global Identifier Slice Selector can be used to associate the packets to the slice aggregate, independent of the MPLS forwarding label that is bound to the destination. LSRs use the MPLS forwarding label to determine the forwarding next-hop(s), and use the Global Identifier Slice Selector field in the packet to infer the specific forwarding treatment that needs to be applied on the packet.

The GISS can be encoded within an MPLS label that is carried in the packet's MPLS label stack. All packets that belong to the same slice aggregate MAY carry the same GISS in the MPLS label stack. It is also possible to have multiple GISS's map to the same slice aggregate. The GISS can be encoded in an MPLS label and may appear in several positions in the MPLS label stack.

2.2.2. Forwarding Label as a Slice Selector

[RFC3031] states in Section 2.1 that: 'Some routers analyze a packet's network layer header not merely to choose the packet's next hop, but also to determine a packet's "precedence" or "class of service"'.

It is possible by assigning a unique MPLS forwarding label to each slice aggregate (FEC) to distinguish the packets forwarded to the same destination but that belong to different slice aggregates. In this case, LSRs can use the top forwarding label to infer both the forwarding action and the forwarding treatment to be invoked on the packets. A similar approach is described in [I-D.ietf-spring-resource-aware-segments] and [I-D.bestbar-teas-ns-packet].

2.3. Time Sensitive Networking

The routers in a network can perform two distinct functions on incoming packets, namely forwarding (where the packet should be sent) and scheduling (when the packet should be sent). Time Sensitive Networking (TSN) and Deterministic Networking provide several mechanisms for scheduling under the assumption that routers are time synchronized. The most effective mechanisms for delay minimization involve per-flow resource allocation.

Segment Routing (SR) is a forwarding paradigm that allows encoding forwarding instructions in the packet in a stack data structure, rather than being programmed into the routers. The SR instructions are contained within a packet in the form of a first-in first-out stack dictating the forwarding decisions of successive routers. Segment routing may be used to choose a path sufficiently short to be capable of providing sufficiently low end- to-end latency but does not influence the queueing of individual packets in each router along that pat

TSN is required for networks transporting time sensitive traffic, that is, packets that are required to be delivered to their final destination by a given time.

2.3.1. Stack-based Methods for Latency Control

One efficient data structure for inserting local deadlines into the headers is a "stack", similar to that used in Segment Routing to carry forwarding instructions. The number of deadline values in the stack equals the number of routers the packet needs to traverse in the network, and each deadline value corresponds to a specific router. The Top-of-Stack (ToS) corresponds to the first router's deadline while the Bottom-of-Stack (BoS) refers to the last's. All local deadlines in the stack are later or equal to the current time (upon which all routers agree), and times closer to the ToS are always earlier or equal to times closer to the BoS.

The ingress router inserts the deadline stack into the packet headers; no other router needs to be aware of the requirements of the time sensitive flows. Hence admitting a new flow only requires updating the information base of the ingress router.

MPLS LSRs that expose the Top of Stack (ToS) label can also inspect the associated "deadline" carried in the packet (either in MPLS stack or after BoS).

2.3.2. Stack Entry Format

A number of different time formats commonly used in networking applications and can be used to encode the local deadlines.

For the forwarding sub-entry we could adopt like SR-MPLS standard 32-bit MPLS labels (which contain a 20-bit label and BoS bit), and thus SR-TSN stack entries could be 64-bits in size comprising a 32-bit MPLS label and the aforementioned nonstandard 32-bit timestamp.

Alternatively, an SR-TSN stack entry could be 96 bits in length comprising a 32-bit MPLS label and either of the standardized 64-bit timestamps.

2.4. NSH Based Service Function Chaining

The Network Service Header (NSH) can be embedded in an Extended Header (EH) to support the Path ID and any metadata that needs to be carried and exchanged between Service Function Forwarders (SFFs).

A reference to the NSH SFC use case is defined in [RFC8596].

2.5. Network Programming

In SR, an ingress node steers a packet through an ordered list of instructions, called "segments". Each one of these instructions represents a function to be called at a specific location in the network. A function is locally defined on the node where it is executed and may range from simply moving forward in the segment list to any complex user-defined behavior.

Network Programming combines Segment Routing (SR) functions to achieve a networking objective that goes beyond mere packet routing.

It may be desirable to encode a pointers to function and its function-arguments within an MPLS packet transport header. For example, in MPLS we can encode the FUNC::ARGs within the label stack or after the bottom of stack to support the equivalent of FUNC::ARG in SRv6 as described in [RFC8986].

2.6. Application Aware Networking (APN)

Application-aware Networking (APN) allows application-aware information (i.e., APN attribute) including APN identification (ID) and/or APN parameters (e.g. network performance requirements) to be encapsulated at network edge devices and carried in packets traversing an APN domain in order to facilitate service provisioning, perform fine-granularity traffic steering and network resource adjustment. To support APN in MPLS networks, mechanisms are needed to hold the APN attribute.

3. Co-existence of Usecases

Two or more of the aforementioned use cases MAY co-exist in the same packet. Some examples of such usecases are described below.

3.1. IOAM with Network Slicing

IOAM may provide key functions with network slicing to help ensure that critical network slice SLOs are being met by the network provider.

In such a case, IOAM is able collect key performance measurement parameters of network slice traffic flows as it traverses the transport network.

This may require, in addition to carrying a specific network slice selector (e.g., GISS), the MPLS network slice packets may have to also carry IOAM ancillary data.

Note that the IOAM ancillary data may have to be modified, and updated on some/all LSRs traversed by the network slice MPLS packets.

3.2. IOAM with Time Sensitive Networking

IOAM operation may also be desirable on MPLS packets that carry time-sensitive related data. Similarly, this may require the presence of multiple ancillary data (whether In-stack or Post-stack ancillary data) to be present in the same MPLS packet.

4. IANA Considerations

This document has no IANA actions.

5. Security Considerations

This document introduces no new security considerations.

6. Acknowledgement

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

7. Contributors

The following individuals contributed to this document:

Kiran Makhijani
Futurewei Technologies
Email: kiranm@futurewei.com

Haoyu Song
Futurewei Technologies
Email: haoyu.song@futurewei.com

Loa Andersson
Bronze Dragon Consulting
Email: loa@pi.nu

8. References

8.1. 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>.

8.2. Informative References

[I-D.bestbar-teas-ns-packet]
Saad, T., Beeram, V. P., Wen, B., Ceccarelli, D., Halpern, J., Peng, S., Chen, R., Liu, X., Contreras, L. M., Rokui, R., and L. Jalil, "Realizing Network Slices in IP/MPLS Networks", Work in Progress, Internet-Draft, draft-bestbar-teas-ns-packet-06, 22 December 2021, <https://www.ietf.org/archive/id/draft-bestbar-teas-ns-packet-06.txt>.

[I-D.decraene-mpls-slid-encoded-entropy-label-id]
Decraene, B., Filsfils, C., Henderickx, W., Saad, T., Beeram, V. P., and L. Jalil, "Using Entropy Label for Network Slice Identification in MPLS networks.", Work in Progress, Internet-Draft, draft-decraene-mpls-slid-encoded-entropy-label-id-02, 6 August 2021, <https://www.ietf.org/archive/id/draft-decraene-mpls-slid-encoded-entropy-label-id-02.txt>.

[I-D.gandhi-mpls-ioam] Gandhi, R., Ali, Z., Brockners, F., Wen, B., Decraene, B., and V. Kozak, "MPLS Data Plane Encapsulation for In-situ OAM Data", Work in Progress, Internet-Draft, draft-gandhi-mpls-ioam-01, 9 September 2021, <https://www.ietf.org/archive/id/draft-gandhi-mpls-ioam-01.txt>.

[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://www.ietf.org/archive/id/draft-gandhi-mpls-ioam-sr-06.txt>.

[I-D.ietf-ippm-ioam-data] Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-data-17, 13 December 2021, <https://www.ietf.org/archive/id/draft-ietf-ippm-ioam-data-17.txt>.

[I-D.ietf-spring-resource-aware-segments] Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Introducing Resource Awareness to SR Segments", Work in Progress, Internet-Draft, draft-ietf-spring-resource-aware-segments-03, 12 July 2021, <https://www.ietf.org/archive/id/draft-ietf-spring-resource-aware-segments-03.txt>.

[I-D.ietf-teas-ietf-network-slices] Farrel, A., Gray, E., Drake, J., Rokui, R., Homma, S., Makhijani, K., Contreras, L. M., and J. Tantsura, "Framework for IETF Network Slices", Work in Progress, Internet-Draft, draft-ietf-teas-ietf-network-slices-05, 25 October 2021, <https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-05.txt>.

[I-D.kompella-mpls-mspl4fa] Kompella, K., Beeram, V. P., Saad, T., and I. Meilik, "Multi-purpose Special Purpose Label for Forwarding Actions", Work in Progress, Internet-Draft, draft-kompella-mpls-mspl4fa-01, 12 July 2021, <https://www.ietf.org/archive/id/draft-kompella-mpls-mspl4fa-01.txt>.

[I-D.li-mpls-enhanced-vpn-vtn-id] Li, Z. and J. Dong, "Carrying Virtual Transport Network Identifier in MPLS Packet", Work in Progress, Internet-Draft, draft-li-mpls-enhanced-vpn-vtn-id-01, 14 April 2021, <https://www.ietf.org/archive/id/draft-li-mpls-enhanced-vpn-vtn-id-01.txt>.

[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>.

[RFC8596] Malis, A., Bryant, S., Halpern, J., and W. Henderickx, "MPLS Transport Encapsulation for the Service Function Chaining (SFC) Network Service Header (NSH)", RFC 8596, DOI 10.17487/RFC8596, June 2019, <https://www.rfc-editor.org/info/rfc8596>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <https://www.rfc-editor.org/info/rfc8986>.

Authors' Addresses

Tarek Saad
Juniper Networks

Email: tsaad@juniper.net

Kiran Makhijani
Futurewei Technologies

Email: kiranm@futurewei.com

Haoyu Song
Futurewei Technologies

Email: haoyu.song@futurewei.com