

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 19, 2020

W. Cheng
H. Li
China Mobile
M. Chen
Huawei
R. Gandhi
Cisco Systems, Inc.
R. Zigler
Broadcom
September 16, 2019

Path Segment in MPLS Based Segment Routing Network
draft-ietf-spring-mpls-path-segment-01

Abstract

A Segment Routing (SR) path is identified by an SR segment list. Only the full segment list identifies the full end-to-end path, and one segment or a partial set of segments from the list cannot identify the SR path and distinguish it from other SR paths that may be partially congruent. But path identification is a pre-requisite for various use-cases such as performance measurement (PM) of an SR path.

In SR for MPLS (SR-MPLS) the segment identifiers are stripped from the packet through label popping as the packet transits the network. This means that when a packet reaches the egress of the SR path, it is not possible to determine on which SR path it traversed the network.

This document defines a new type of segment that is referred to as Path Segment, which is used to identify an SR path in an SR-MPLS network. When used, it is inserted at the ingress node of the SR path and immediately follows the last segment identifier in the segment list of the SR path. The Path Segment will not be popped off until it reaches the egress node of the SR path.

A Path Segment can be used by the egress node to implement path identification hence to support various use-cases including SR path PM, end-to-end 1+1 SR path protection, and bidirectional SR paths correlation.

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 March 19, 2020.

Copyright Notice

Copyright (c) 2019 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	3
1.1.	Requirements Language	3
1.2.	Abbreviations	3
2.	Path Segment	4
3.	Nesting of Path Segments	5
4.	Path Segment Allocation	6
5.	Path Segment for PM	7
6.	Path Segment for Bi-directional SR Path	7
7.	Path Segment for End-to-end Path Protection	7
8.	IANA Considerations	8
9.	Security Considerations	8
10.	Contributors	8
11.	Acknowledgements	8
12.	References	8
12.1.	Normative References	9
12.2.	Informative References	9
	Authors' Addresses	10

1. Introduction

Segment Routing (SR) [[RFC8402](#)] is a source routed forwarding method that allows to directly encode forwarding instructions (called segments) in each packet, hence it enables steering traffic through a network without the per-flow states maintained on the transit nodes. Segment Routing can be instantiated on an MPLS data plane or an IPv6 data plane. The former is called SR-MPLS [[I-D.ietf-spring-segment-routing-mpls](#)], the latter is called SRv6 [[RFC8402](#)]. SR-MPLS leverages the MPLS label stack to construct SR paths.

In an SR-MPLS network, when a packet is transmitted along an SR path, the labels in the MPLS label stack will be swapped or popped. So that no label or only the last label may be left in the MPLS label stack when the packet reaches the egress node. Thus, the egress node cannot determine along which SR path the packet came.

However, to support use cases like end-to-end 1+1 path protection (Live-Live case), bidirectional path correlation or performance measurement (PM), the ability to implement path identification is a pre-requisite. More detail can be found in [Section 5](#), 6, and 7.

Therefore, this document introduces a new segment that is referred to as the Path Segment. A Path Segment is defined to uniquely identify an SR path in an SR-MPLS network in the context of the egress node. It is normally used by egress nodes for path identification or correlation.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP14](#) [[RFC2119](#)][[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

1.2. Abbreviations

DM: Delay Measurement.

LM: Loss Measurement.

MPLS: Multiprotocol Label Switching.

PM: Performance Measurement.

PSID: Path Segment ID.

SID: Segment ID.

SL: Segment List.

SR: Segment Routing.

SR-MPLS: Segment Routing instantiated on MPLS data plane.

SRv6: Segment Routing instantiated on IPv6 data plane

2. Path Segment

A Path Segment is a single label that is assigned from the Segment Routing Local Block (SRLB) or Segment Routing Global Block (SRGB) of the egress node of an SR path. It means that the Path Segment is unique in the context of the egress node of the SR path. When a Path Segment is used, the Path Segment MUST be inserted at the ingress node and MUST immediately follow the last label of the SR path. The Path Segment may be used to identify an SR-MPLS Policy, its Candidate-Path (CP), or a SID List (SL) [[I-D.ietf-spring-segment-routing-policy](#)] terminating on an egress node depending on the use-case.

The value of the TTL field in the MPLS label stack entry containing the Path Segment MUST be set to the same value as the TTL of the last label stack entry for the last segment in the SR path. If the Path Segment is the bottom label, the S bit MUST be set.

Normally, the intermediate nodes will not see the Path Segment label and do not know how to process it. A Path Segment presenting to an intermediate node is an error condition.

The egress node MUST pop the Path Segment. The egress node MAY use the Path Segment for further processing. For example, when performance measurement is enabled on the SR path, it can trigger packet counting or timestamping.

The label stack with Path Segment is as below (Figure 1):



Figure 1: Label Stack with Path Segment

Where:

- o The Labels 1 to n are the segment label stack used to direct how to steer the packets along the SR path.
- o The Path Segment identifies the SR path in the context of the egress node of the SR path.

3. Nesting of Path Segments

Binding SID (BSID) [[RFC8402](#)] can be used for SID list compression. With BSID, an end-to-end SR path can be split into several sub-paths, each sub-path is identified by a BSID. Then an end-to-end SR path can be identified by a list of BSIDs, therefore, it can provide better scalability.

BSID and Path SID (PSID) can be combined to achieve both sub-path and end-to-end path monitoring. A reference model for such a combination in (Figure 2) shows an end-to-end path (A->D) that spans three domains (Access, Aggregation and Core domain) and consists of three sub-paths, one in each sub-domain (sub-path (A->B), sub-path (B->C) and sub-path (C->D)). Each sub-path is allocated a BSID. For nesting the sub-paths, each sub-path is allocated a PSID. Then, the SID list of the end-to-end path can be expressed as <BSID1, BSID2, ..., BSIDn, e-PSID>, where the e-PSID is the PSID of the end-to-end path. The SID list of a sub-path can be expressed as <SID1, SID2, ...SIDn, s-PSID>, where the s-PSID is the PSID of the sub-path.

Figure 2 shows the details of the label stacks when PSID and BSID are used to support both sub-path and end-to-end path monitoring in a multi-domain scenario.




Figure 2: Nesting of Path Segments

4. Path Segment Allocation

Several ways can be used to allocate the Path Segment.

One way is to set up a communication channel (e.g., MPLS Generic Associated Channel (G-ACh)) between the ingress node and the egress node, and the ingress node of the SR path can directly send a request to the egress node to ask for a Path Segment.

Another way is to leverage a centralized controller (e.g., PCE, SDN controller) to assign the Path Segment. PCEP based Path Segment allocation is defined in [\[I-D.li-pce-sr-path-segment\]](#), and SR-policy based path segment allocation is defined in [\[I-D.li-idr-sr-policy-path-segment-distribution\]](#).

A Path Segment can be used in the case of PHP, where some labels MAY be popped off at the penultimate hop of an SR path, but the Path Segment MUST NOT be popped off until it reaches at the egress LSR. In the case of a Path Segment that is assigned by a centralized controller, the controller must make sure (e.g., by some capability

discovery mechanisms) that the egress LSR knows and can process the Path Segment.

5. Path Segment for PM

As defined in [[RFC7799](#)], performance measurement can be classified into Active, Passive and Hybrid measurement. For Passive measurement, path identification at the measuring points is the prerequisite. Path segment can be used by the measuring points (e.g., the ingress/egress nodes of an SR path) or a centralized controller to correlate the packets counts/timestamps that are from the ingress and egress nodes to a specific SR path, then packet loss/delay can be calculated.

Performance Delay Measurement (DM) and Loss Measurements (LM) in SR networks with MPLS data plane can be found in [[I-D.gandhi-spring-sr-mpls-pm](#)] and [[I-D.gandhi-spring-udp-pm](#)].

6. Path Segment for Bi-directional SR Path

With the current SR architecture, an SR path is a unidirectional path. In some scenarios, for example, mobile backhaul transport network, there are requirements to support bidirectional paths, and the path is normally treated as a single entity and both directions of the path have the same fate, for example, failure in one direction will result in switching at both directions.

MPLS supports this by introducing the concepts of co-routed bidirectional LSP and associated bidirectional LSP. With SR, to support bidirectional paths, a straightforward way is to bind two unidirectional SR paths to a single bidirectional path. Path Segments can be used to correlate the two unidirectional SR paths at both ends of the paths.

[[I-D.li-pce-sr-bidir-path](#)] defines how to use PCEP and Path segment to initiate a bidirectional SR path, and [[I-D.li-idr-sr-policy-path-segment-distribution](#)] defines how to use SR policy and Path segment to initiate a bidirectional SR path.

7. Path Segment for End-to-end Path Protection

For end-to-end 1+1 path protection (i.e., Live-Live case), the egress node of an SR path needs to know the set of paths that constitute the primary and the secondaries, in order to select the primary packet for onward transmission, and to discard the packets from the secondaries.

To do this, each path needs a path identifier that is unique at the egress node. Depending on the design, this is a single unique path segment label chosen by the egress PE.

There then needs to be a method of binding this path identifiers into equivalence groups such that the egress PE can determine the set of packets that represent a single path and its secondary.

It is obvious that this group can be instantiated in the network by an SDN controller and that the Path Segment achieves the desired function.

8. IANA Considerations

This document does not require any IANA actions.

9. Security Considerations

This document does not introduce additional security requirements and mechanisms other than the ones described in [[RFC8402](#)].

10. Contributors

The following individuals also contribute to this document.

- o Cheng Li, Huawei
- o Lei Wang, China Mobile
- o Shuangping Zhan, ZTE
- o Greg

11. Acknowledgements

The authors would like to thank Adrian Farrel, Stewart Bryant, Alexander Vainshtein, Andrew G. Malis and Loa Andersson for their review, suggestions and comments to this document.

The authors would like to acknowledge the contribution from Alexander Vainshtein on "Nesting of Path Segments".

12. References

12.1. Normative References

- [I-D.ietf-spring-segment-routing-mpls]
Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with MPLS data plane", [draft-ietf-spring-segment-routing-mpls-22](#) (work in progress), May 2019.
- [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>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

12.2. Informative References

- [I-D.gandhi-spring-sr-mpls-pm]
Gandhi, R., Filsfils, C., daniel.voyer@bell.ca, d., Salsano, S., Ventre, P., and M. Chen, "Performance Measurement in Segment Routing Networks with MPLS Data Plane", [draft-gandhi-spring-sr-mpls-pm-03](#) (work in progress), September 2018.
- [I-D.gandhi-spring-udp-pm]
Gandhi, R., Filsfils, C., daniel.voyer@bell.ca, d., Salsano, S., Ventre, P., and M. Chen, "UDP Path for In-band Performance Measurement for Segment Routing Networks", [draft-gandhi-spring-udp-pm-02](#) (work in progress), September 2018.
- [I-D.ietf-6man-segment-routing-header]
Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header (SRH)", [draft-ietf-6man-segment-routing-header-22](#) (work in progress), August 2019.

[I-D.ietf-spring-segment-routing-policy]

Filsfils, C., Sivabalan, S., daniel.voyer@bell.ca, d., bogdanov@google.com, b., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-03](#) (work in progress), May 2019.

[I-D.li-idr-sr-policy-path-segment-distribution]

Li, C., Chen, M., Dong, J., and Z. Li, "Segment Routing Policies for Path Segment and Bidirectional Path", [draft-li-idr-sr-policy-path-segment-distribution-01](#) (work in progress), October 2018.

[I-D.li-pce-sr-bidir-path]

Li, C., Chen, M., Cheng, W., Li, Z., Dong, J., Gandhi, R., and Q. Xiong, "PCEP Extensions for Associated Bidirectional Segment Routing (SR) Paths", [draft-li-pce-sr-bidir-path-06](#) (work in progress), August 2019.

[I-D.li-pce-sr-path-segment]

Li, C., Chen, M., Cheng, W., Dong, J., Li, Z., Gandhi, R., and Q. Xiong, "Path Computation Element Communication Protocol (PCEP) Extension for Path Segment in Segment Routing (SR)", [draft-li-pce-sr-path-segment-08](#) (work in progress), August 2019.

[RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", [RFC 7799](#), DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.

Authors' Addresses

Weiqiang Cheng
China Mobile

Email: chengweiqiang@chinamobile.com

Han Li
China Mobile

Email: lihan@chinamobile.com

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

Rakesh Gandhi
Cisco Systems, Inc.
Canada

Email: rgandhi@cisco.com

Royi Zigler
Broadcom

Email: royi.zigler@broadcom.com