

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 3, 2018

W. Cheng
L. Wang
H. Li
China Mobile
M. Chen
Huawei
R. Zigler
Broadcom
S. Zhan
ZTE
October 30, 2017

Path Segment in MPLS Based Segment Routing Network
draft-cheng-spring-mpls-path-segment-00

Abstract

An SR path is identified by an SR segment list, one or partial segments of the list cannot uniquely identify the SR path.

This document introduces the concept of Path Segment that is used to identify an SR path. When used, it is inserted at the ingress node of the SR path and immediately follows the last segment of the SR path. The Path Segment will not be popped off until it reaches the egress of the SR path, it can be used by the egress node to implement end-2-end SR path protection or performance measurement (PM) of an SR path.

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 May 3, 2018.

Copyright Notice

Copyright (c) 2017 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
2. Path Segment	3
2.1. One Label Solution	3
2.1.1. Path Segment Assignment	4
2.2. Two Labels Solution	5
3. Path Segment Application	7
3.1. Performance Measurement	7
3.2. End-2-end Path Protection	7
3.3. Bi-directional SR Tunnel	8
4. IANA Considerations	8
5. Security Considerations	8
6. Contributors	8
7. Acknowledgements	8
8. References	8
8.1. Normative References	8
8.2. Informative References	9
Authors' Addresses	9

1. Introduction

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

[I-D.ietf-spring-segment-routing-mpls], the latter is called SRv6 [I-D.ietf-6man-segment-routing-header]. SR-MPLS leverages the MPLS label stack to construct SR path, and SRv6 uses the Segment Routing Header to construct SR path.

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 from which ingress node or SR path the packet comes.

However, to support use cases like end-2-end 1+1 path protection, bidirectional path correlation or performance measurement(PM), the ability to implement path identification is the pre-condition.

Therefore, this document introduces a new segment that is referred to as Path Segment. A Path Segment is defined to unique identify an SR path in a specific context. (e.g., in the context of the egress node or ingress node of an SR path, or within an SR domain). It is normally used by egress nodes for path identification or correlation. Path Segment can only apply to SR-MPLS.

2. Path Segment

This document introduces two options for SR path identification: one label solution and two labels solution.

[Editor notes: it is supposed that the WG will discuss and decide which one is the better solution.]

2.1. One Label Solution

The 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 SR paths. When Path Segment is used, a Path label MUST be inserted at the ingress node and MUST immediately follow the last label of the SR path.

If the Path label is the bottom label, the S bit MUST be set. The value of the TC field MUST be set to the same value as the last segment label of the SR path. The value of the TTL field MUST be set to the same value of the last segment label of the SR path.

Normally, the intermediates node will not see the Path Segment label and do not know how to process it even if they see it. A Path Segment label presenting to an intermediate node is error situation.

The egress node MUST pop the Path label and deliver it to relevant components for further processing.

The label stack with Path Segment is as below (Figure1):




Figure 1: Label Stack with Path Segment

Where:

- o The Label 1-n are the segment labels that are used to direct how to steer the packets along the SR path.
- o The Path Label identifies the SR path in the context of the egress node of the SR path.

2.1.1. Path Segment Assignment

Several ways can be used to assign the Path Segment. One way is that the Path Segment label is directly assigned by the egress node of an SR path. Where the ingress node of the SR path can directly send a request to the egress node to ask for a Path label. With this way, it needs to set up a communication channel between the ingress node and the egress node. New protocols or extensions to existing protocol may be required.

Another candidate way is to leverage a centralized controller (e.g., PCE) to assign the Path label. The ingress node sends a request to the PCE to compute a SR path and indicate that a Path label is desired. The PCE will compute the path as required. Once the path computed, the PCE will send a request (with computed path and relevant information) to the egress node to ask for a Path label for

the SR path. The egress node will allocate a label to the SR path and build mapping relationship between the label and the path. A reply will be sent back to the PCE, the PCE will send a reply to the ingress node about the path information and the corresponding Path Segment label.

With either way or the variations, the final purpose is to assign a label from the egress node's label space, hence a single label is enough for path identification. Then the ingress node can put the Path Segment label into the label stack when needed, and the egress node can use that Path Segment to implement relevant functionalities.

2.2. Two Labels Solution

Two segments (Source segment and Path segment) are used to identify an SR path. The Source segment is a global node segment, it can uniquely identify a node within an SR domain. It MUST NOT be used for forwarding and indicates that a Path segment immediately follows. The Path segment is a local segment generated at the ingress node to identify an SR path. The combination of Source segment and Path segment can uniquely identify an SR Path with an SR domain.

A node that enables Path segment function will be assigned two node segments. One is for forwarding just as defined in [I-D.ietf-spring-segment-routing], the other is for source identification. The corresponding label of the Source Segment is indexed in the SRGB (or in a of the node to which the Source Segment will be presented).

The Path segment label is a local label that is assigned to an SR path at the ingress node.

The label stack with Source and Path segments is as below (Figure 2):




Figure 2: Label Stack with Source and Path Segments

Where:

- o The Label 1-n are the segment labels that are used to direct how to steer the packets along an SR path, and the "label n" is the last label of the SR path or the label that directs forwarding packets to the node to which the Source Segment will be presented.
- o The Source Label identifies the source of the SR path. The value of the TC and TTL fields of the Source Label MUST be set to the same values as the label (e.g., the Label n) it follows.
- o The Path Label identifies the SR path in the context of source node. If the Path label is the bottom label, the S bit MUST be set. The value of the TC and TTL fields SHOULD be set to the same values as the Source label.

The Source and Path label MUST be inserted at the ingress node of an SR path. And they MUST immediately follow the label that directs forwarding packets to the node (e.g., the egress or an intermediate node) to which the Source Segment (as the stack top label) and Path Segment are presented.

If a node receives a packet with an unknown Source Label, the packet MUST be discarded and an error SHOULD be reported.

The Source label and Path label MUST be popped at the node who receives a packet with the Source label as the stack top label.

3. Path Segment Application

3.1. Performance Measurement

To measure the packet loss and delay of the real traffic of an SR path, one fundamental condition is path identification at the measuring points. For an SR path, the ingress node have the complete information of the path, it can use those information for packet counting and/or timestamping. At the egress node, since the Path Segment label (or combination with Source label) can be used to identify the path, path based packet counting and/or timestamping can be implemented as well. Then combined with the mechanisms defined [RFC6374], end-2-end packet loss and/or delay measurement of an SR path can be achieved.

Measuring at intermediate nodes needs more consideration, it will be added in the next version.

3.2. End-2-end Path Protection

For end-2-end 1+1 path protection, the egress node of an path needs to know the set of paths that constitute the primary and the backup(s), in order to select the primary packet for onward transmission, and to discard the packets from the backups.

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

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 backup.

It is obvious that this group can be instantiated in the network by an SDN controller.

In a network that is using a distributed control plane the approach will depend on the control protocol used, but the essence of the solution is that which ever PE is responsible for creating the group advertises then as a group of equivalent paths. Whether one of these is advertised as primary and the others as secondary will or all are advertised as of equal status will depend on the details of the underlying protection mechanism.

3.3. Bi-directional SR Tunnel

With the current SR architecture, an SR path is an unidirectional path. In some scenarios, for example, mobile backhaul transport network, there are requirements to support bi-directional path, and the path is normally treated as a single entity and both directions of the path have 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 bi-directional LSP. With SR, to support bidirectional path, a straightforward way is to bind two unidirectional SR paths to a single bi-directional path. Path segments can be used to correlate the two unidirectional SR paths at both ends of the paths.

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

6. Contributors

The following individuals also contribute to this document.

- o Shuangping Zhan, ZTE
- o Cheng Li, Huawei

7. Acknowledgements

The authors would like to thank Stewart Bryant for his review, suggestion and comments to this document.

8. References

8.1. Normative References

- [I-D.ietf-6man-segment-routing-header]
Previdi, S., Filsfils, C., Raza, K., Leddy, J., Field, B.,
daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d.,
Matsushima, S., Leung, I., Linkova, J., Aries, E., Kosugi,
T., Vyncke, E., Lebrun, D., Steinberg, D., and R. Raszuk,
"IPv6 Segment Routing Header (SRH)", draft-ietf-6man-
segment-routing-header-07 (work in progress), July 2017.
- [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>>.

8.2. Informative References

- [I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B.,
Litkowski, S., and R. Shakir, "Segment Routing
Architecture", draft-ietf-spring-segment-routing-13 (work
in progress), October 2017.
- [I-D.ietf-spring-segment-routing-mpls]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B.,
Litkowski, S., and R. Shakir, "Segment Routing with MPLS
data plane", draft-ietf-spring-segment-routing-mpls-10
(work in progress), June 2017.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay
Measurement for MPLS Networks", RFC 6374,
DOI 10.17487/RFC6374, September 2011,
<<https://www.rfc-editor.org/info/rfc6374>>.

Authors' Addresses

Weiqliang Cheng
China Mobile

Email: chengweiqliang@chinamobile.com

Lei Wang
China Mobile

Email: wangleiyj@chinamobile.com

Han Li
China Mobile

Email: lihan@chinamobile.com

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

Royi Zigler
Broadcom

Email: royi.zigler@broadcom.com

Shuangping Zhan
ZTE

Email: zhan.shuangping@zte.com.cn