

SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 29, 2021

C. Li
Huawei Technologies
W. Cheng
China Mobile
M. Chen
D. Dhody
Huawei Technologies
R. Gandhi
Cisco Systems, Inc.
November 25, 2020

Path Segment for SRv6 (Segment Routing in IPv6)
draft-ietf-spring-srv6-path-segment-00

Abstract

Segment Routing (SR) allows for a flexible definition of end-to-end paths by encoding an ordered list of instructions, called "segments". The SR architecture can be implemented over an MPLS data plane as well as an IPv6 data plane.

Currently, Path Segment has been defined to identify an SR path in SR-MPLS networks, and is used for various use-cases such as end-to-end SR Path Protection and Performance Measurement (PM) of an SR path. This document defines the Path Segment to identify an SRv6 path in an IPv6 network.

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 29, 2021.

Copyright Notice

Copyright (c) 2020 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
1.1.	Requirements Language	4
1.2.	Terminology	4
2.	Use Cases for SRv6 Path Segment	4
3.	SRv6 Path Segment	5
3.1.	Format of an SRv6 Path Segment	5
3.1.1.	SRv6 Path Segment: Locator and Local ID	5
3.1.2.	SRv6 Path Segment: Global ID	6
4.	SRv6 Path Segment Allocation	6
5.	Processing of SRv6 Path Segment	7
6.	IANA Considerations	7
7.	Security Considerations	8
8.	Contributors	8
9.	Acknowledgements	8
10.	References	8
10.1.	Normative References	8
10.2.	Informative References	9
	Authors' Addresses	10

[1.](#) Introduction

Segment routing (SR) [[RFC8402](#)] is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node by inserting an ordered list of instructions, called segments.

When segment routing is deployed on an MPLS data plane, called SR-MPLS [[I-D.ietf-spring-segment-routing-mpls](#)], a segment identifier (SID) is present as an MPLS label. When segment routing is deployed on an IPv6 data plane, a SID is presented as a 128-bit value, and it can be an IPv6 address of a local interface but it does not have to

be. To support SR in an IPv6 network, a Segment Routing Header (SRH) [[RFC8754](#)] is used.

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 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 can not determine from which ingress node or SR path the packet came from. Therefore, to identify an SR-MPLS path, a Path Segment is defined in [[I-D.ietf-spring-mpls-path-segment](#)].

Likewise, a path needs to be identified in an SRv6 network for several use cases such as binding bidirectional paths [[I-D.ietf-pce-sr-bidir-path](#)] and end-to-end performance measurement [[I-D.gandhi-spring-udp-pm](#)].

An SRv6 path MAY be identified by the content of a segment list. However, the segment list may not be a good key, since the length of a segment list is flexible according to the number of required SIDs. Also, the length of a segment list may be too long to be a key when it contains many SIDs. For instance, if packet A uses an SRH with 3 SIDs while Packet B uses an SRH with 10 SIDs, the key to identify these two paths will be a 384-bits value and a 1280-bits value, respectively. Further, an SRv6 path cannot be identified by the information carried by the SRH in reduced mode [[RFC8754](#)] as the first SID is not present.

Furthermore, different SRv6 policies may use the same segment list for different candidate paths, so the traffic of different SRv6 policies are merged, resulting in the inability to measure the performance of the specific path.

To solve the above issues, this document defines a new SRv6 segment called "SRv6 Path Segment", which is a 128-bits value, to identify an SRv6 path.

When the SRv6 Path Segment is used in reduced mode SRH [[RFC8754](#)], the entire path information is indicated by the Path Segment, and the performance will be better than using the entire segment list as the path identifier, while the overhead is equivalent to the SRH in normal mode. Furthermore, with SRv6 Path Segment, each SRv6 candidate path can be identified and measured, even when they use the same segment list.

An SRv6 Path Segment MUST NOT be copied to the IPv6 destination address, so it is not routable.

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 [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

MPLS: Multiprotocol Label Switching.

PM: Performance Measurement.

SID: Segment ID.

SR: Segment Routing.

SR-MPLS: Segment Routing with MPLS data plane.

SRH: Segment Routing Header.

PSID: Path Segment Identifier.

PSP: Penultimate Segment Popping.

Further, this document makes use of the terms defined in [[RFC8402](#)] and [[I-D.ietf-spring-srv6-network-programming](#)].

2. Use Cases for SRv6 Path Segment

Similar to SR-MPLS Path Segment [[I-D.ietf-spring-mpls-path-segment](#)], SRv6 Path Segment may also be used to identify an SRv6 Path in some use cases:

- o Performance Measurement: For Passive measurement [[RFC7799](#)], path identification at the measuring points is the pre-requisite [[I-D.ietf-spring-mpls-path-segment](#)]. SRv6 Path segment can be used by the measuring points (e.g., the ingress/egress nodes of an SRv6 path) or a centralized controller to correlate the packets counts/timestamps, then packet loss/delay can be calculated.
- o Bi-directional SRv6 Path Association: In some scenarios, such as mobile backhaul transport networks, there are requirements to support bidirectional paths. Like SR-MPLS [[I-D.ietf-spring-mpls-path-segment](#)], to support bidirectional SRv6 paths, a straightforward way is to bind two unidirectional SRv6 paths to a single bidirectional path. SRv6 Path segments can be

used to correlate the two unidirectional SRv6 paths at both ends of the path. [[I-D.ietf-pce-sr-bidir-path](#)] defines how to use PCEP and Path Segment to initiate a bidirectional SR path.

- o End-to-end Path Protection: For end-to-end 1+1 path protection (i.e., Live-Live case), the egress node of an SRv6 path needs to know the set of paths that constitute the primary and the secondary(s), to select the primary packet for onward transmission, and to discard the packets from the secondary(s), so each SRv6 path needs a unique path identifier at the egress node, which can be an SRv6 Path Segment.

3. SRv6 Path Segment

As defined in [[I-D.ietf-spring-srv6-network-programming](#)], an SRv6 segment is a 128-bit value.

To identify an SRv6 path, this document defines a new segment called SRv6 Path Segment.

The SRv6 Path Segment MUST appear only once in a segment list, and it MUST appear as the last entry in the segment list. To indicate the SRv6 Path Segment, an SRH.P-flag is defined in [[I-D.li-6man-srv6-path-segment-encap](#)].

Depending on the use case, an SRv6 Path Segment identifies:

- o an SRv6 path within an SRv6 domain
- o an SRv6 Policy
- o a Candidate-path or a SID-List in a SRv6 Policy
[[I-D.ietf-spring-segment-routing-policy](#)]

Note that, based on the use-case, a SRv6 Path Segment can be used for different SID-Lists within an SR Policy.

3.1. Format of an SRv6 Path Segment

This document proposes two types of SRv6 Path Segment format.

3.1.1. SRv6 Path Segment: Locator and Local ID

As per [[I-D.ietf-spring-srv6-network-programming](#)], an SRv6 segment is a 128-bit value, which can be represented as LOC:FUNCT, where LOC is the L most significant bits and FUNCT is the 128-L least significant bits. L is called the locator length and is flexible. Each network operator is free to use the locator length it chooses. Most often

the LOC part of the SID is routable and leads to the node which instantiates that SID. The FUNCT part of the SID is an opaque identification of a local function bound to the SID. The FUNCT value zero is invalid.

SRv6 Path Segment can follow the format, where the LOC part identifies the egress node that allocates the Path Segment, and the FUNCT part is a unique local ID to identify an SRv6 Path and its endpoint behavior.

The Function Type of an SRv6 Path Segment is END.PSID (End Function with Path Segment Identifier).



Figure 2. PSID in Format LOC:FUNCT

3.1.2. SRv6 Path Segment: Global ID

An SRv6 Path Segment ID can be a Global ID, and its format depends on the use case.

The SRv6 Path Segment will not be copied to the IPv6 Destination Address, so the SRv6 Path Segment ID can be allocated from an independent 128-bits ID Space. In this case, a new table should be maintained at the node for SRv6 Path Segment.




Figure 3. A Global ID as an PSID

4. SRv6 Path Segment Allocation

A Path Segment is a local segment allocated by an egress node. A Path Segment can be allocated through several ways, such as CLI, BGP [[I-D.ietf-idr-sr-policy-path-segment](#)], PCEP [[I-D.ietf-pce-sr-path-segment](#)] or other ways. The mechanisms through which a Path Segment is allocated are out of scope of this document.

When a Path Segment is allocated by the egress, it MUST be distributed to the ingress node of the path that identified by the path segment. In this case, only the egress will process the Path Segment, and other nodes specified by SIDs in the segment list do not know how to process the Path Segment.

Depending on the use case, a Path Segment may be distributed to the SRv6 nodes along the SRv6 path. In this case, the SRv6 nodes that learned the Path Segment may process the Path Segment depending on the use case.

5. Processing of SRv6 Path Segment

When the SRv6 Path Segment is used, the following rules apply:

- o The SRv6 Path Segment MUST appear only once in a segment list, and it MUST appear as the last entry. Only the one that appears as the last entry in the SID list will be processed. An SRv6 Path Segment that appears at any other location in the SID list will be treated as an error.
- o When an SRv6 Path Segment is inserted, the SL MUST be initiated to be less than the value of Last Entry, and will not point to SRv6 Path Segment. For instance, when the Last entry is 4, the SID List[4] is the SRv6 Path Segment, so the SL MUST be set to 3 or other numbers less than Last entry.
- o The SRv6 Path Segment MUST NOT be copied to the IPv6 destination address.
- o Penultimate Segment Popping (PSP, as defined in [\[I-D.ietf-spring-srv6-network-programming\]](#)) MUST be disabled.
- o The ingress needs to set the P-bit when an SRv6 Path Segment is inserted in the SID List. Nodes that support SRv6 Path Segment processing will inspect the last entry to process SRv6 Path Segment when the P-bit is set. When the P-bit is unset, the nodes will not inspect the last entry.
- o The specific SRv6 Path Segment processing depends on use cases, and it is out of scope of this document.

6. IANA Considerations

This I-D requests the IANA to allocate, within the "SRv6 Endpoint Behaviors" sub-registry belonging to the top-level "Segment-routing with IPv6 data plane (SRv6) Parameters" registry, the following allocations:

Value	Description	Reference

TBA1	End.PSID - SRv6 Path Segment	[This.ID]

7. Security Considerations

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

8. Contributors

Zhenbin Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

9. Acknowledgements

The authors would like to thank Stefano Previdi and Zafar Ali for their valuable comments and suggestions.

10. References

10.1. Normative References

[I-D.ietf-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., Voyer, D.,
Matsushima, S., and Z. Li, "SRv6 Network Programming",
[draft-ietf-spring-srv6-network-programming-24](#) (work in
progress), October 2020.

- [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>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", [RFC 8754](#), DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

[10.2](#). Informative References

- [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-idr-sr-policy-path-segment]
Li, C., Li, Z., Chen, H., Cheng, W., and K. Talaulikar, "SR Policy Extensions for Path Segment and Bidirectional Path", [draft-ietf-idr-sr-policy-path-segment-02](#) (work in progress), November 2020.
- [I-D.ietf-pce-sr-bidir-path]
Li, C., Chen, M., Cheng, W., Gandhi, R., and Q. Xiong, "PCEP Extensions for Associated Bidirectional Segment Routing (SR) Paths", [draft-ietf-pce-sr-bidir-path-03](#) (work in progress), September 2020.
- [I-D.ietf-pce-sr-path-segment]
Li, C., Chen, M., Cheng, W., Gandhi, R., and Q. Xiong, "Path Computation Element Communication Protocol (PCEP) Extension for Path Segment in Segment Routing (SR)", [draft-ietf-pce-sr-path-segment-02](#) (work in progress), November 2020.

[I-D.ietf-spring-mpls-path-segment]

Cheng, W., Li, H., Chen, M., Gandhi, R., and R. Zigler,
"Path Segment in MPLS Based Segment Routing Network",
[draft-ietf-spring-mpls-path-segment-03](#) (work in progress),
September 2020.

[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.

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

Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and
P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-09](#) (work in progress),
November 2020.

[I-D.li-6man-srv6-path-segment-encap]

Li, C., Cheng, W., Li, Z., and D. Dhody, "Encapsulation of
Path Segment in SRv6", [draft-li-6man-srv6-path-segment-encap-04](#) (work in progress), October 2020.

[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

Cheng Li
Huawei Technologies

Email: c.l@huawei.com

Weiqiang Cheng
China Mobile

Email: chengweiqiang@chinamobile.com

Mach(Guoyi) Chen
Huawei Technologies

Email: mach.chen@huawei.com

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India

Email: dhruv.ietf@gmail.com

Rakesh Gandhi
Cisco Systems, Inc.
Canada

Email: rgandhi@cisco.com

