

SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: 5 October 2025

C. Li
Huawei Technologies
W. Cheng
China Mobile
M. Chen
D. Dhody
Huawei Technologies
Y. Zhu
China Telecom
3 April 2025

Path Segment Identifier (PSID) in SRv6 (Segment Routing in IPv6)
draft-ietf-spring-srv6-path-segment-12

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 Identifier (PSID) 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 PSID 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 5 October 2025.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	4
1.2. Terminology	4
2. Use Cases for SRv6 PSID	4
3. SRv6 Path Segment Identifier	5
3.1. Structure of an SRv6 PSID	6
3.1.1. Reusing Existing SID Structure	6
4. Encoding of an SRv6 PSID	7
4.1. SRH.P-flag	7
5. SRv6 PSID Allocation	9
6. Processing of SRv6 PSID	9
7. IANA Considerations	10
8. Security Considerations	10
9. Contributors	11
10. Acknowledgements	11
11. References	12
11.1. Normative References	12
11.2. Informative References	12
Authors' Addresses	13

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 [RFC8660], 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 SR, a path needs to be identified for several use cases such as binding bidirectional paths [I-D.ietf-pce-sr-bidir-path] and end-to-end performance measurement [I-D.ietf-spring-stamp-srpm].

Additionally, 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 cannot determine from which ingress node or SR path the packet came. To identify an SR-MPLS path, a Path Segment Identifier is defined in [RFC9545].

An SRv6 path could be identified by the content of a segment list. However, the segment list is not be a good key identifier, 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 in the SRH.

In the cases that different SRv6 policies use the same segment list in different candidate paths, the traffic of different SRv6 policies are merged on the egress segment node because the segment list in the SRH cannot distinguish the path in different SR policy, resulting in the inability to measure the performance of a specific path within its SRv6 policy. However, an SRv6 policy may need to measure a segment list in its own candidate path, where the packets statistics associated with this segment list is independent with other segment lists in other SRv6 policies even if the same segment list is used. Without a Path ID to specify the path will cause the statistic of the traffic from multiple paths using the same segment list under different SRv6 policies merge into an aggregated result on the egress endpoint node.

To solve the above issues, this document defines a new SRv6 segment called the "SRv6 Path Segment Identifier (PSID)", which in total is an 128-bits value, to identify an SRv6 path. Using a 128-bit Path ID has the better processing performance than using a flexible length of Segment list as a path ID. Using an SRv6 PSID is used in reduced mode SRH [RFC8754] can solve the problem of cannot identifying a segment list by the reduced segment list, while the overhead is equivalent to the SRH in normal mode. Furthermore, by using SRv6 PSIDs, any SRv6 path within an SRv6 policy can be identified and measured, even when they use the same segment list.

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

PM: Performance Measurement.

SID: Segment Identifier.

SR: Segment Routing.

SR-MPLS: Segment Routing with MPLS data plane.

SRH: Segment Routing Header.

SR path: A path described by a segment list [RFC9545].

SRv6 path: A path described by an SRv6 segment list.

PSID: Path Segment Identifier.

PSP: Penultimate Segment Popping.

Further, this document makes use of the terms defined in [RFC8402] and [RFC8986].

2. Use Cases for SRv6 PSID

Similar to SR-MPLS PSID [RFC9545], SRv6 PSID may also be used to identify an SRv6 Path in some use cases:

- * Performance Measurement: For Passive measurement [RFC7799], path identification at the measuring points is the pre-requisite [RFC9545]. SRv6 PSIDs 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.
- * Bi-directional SRv6 Path Association: In some scenarios, such as mobile backhaul transport networks, there are requirements to support bidirectional paths. Like SR-MPLS [RFC9545], to support bidirectional SRv6 paths, a straightforward way is to bind two unidirectional SRv6 paths to a single bidirectional path. SRv6

PSIDs 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 PSID to initiate a bidirectional SR path.

- * 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 PSID.

3. SRv6 Path Segment Identifier

As defined in [RFC8986], an SRv6 segment identifier is a 128-bit value.

To identify an SRv6 path, this document defines a new segment called SRv6 Path Segment, and the Path Segment Identifier (PSID) is a 128-bit value. An SRv6 PSID MUST NOT be used for routing so it MUST NOT be copied to the IPv6 destination address. [RFC8754] states that the SR segment endpoint node creates Forwarding Information Base (FIB) entries for its local SIDs (without constraining the details of implementation). In order to provide a new independent 128-bit ID space for PSID, the PSID is required to be stored separating from the other SIDs (for example in a different table from the FIB).

A PSID is used for identify a SRv6 path represented by a segment list. On the egress node, the statistics of this path will be identified by its PSID. In normal cases, each segment list will have a its own PSID with different value. However, depending on the use case, different segment list might use the same value PSID, causing the statistics of these SRv6 path are merged on the egress node. For example, a use case may use the same PSID for some or all the segment list under a Candidate path, or even some or all Candidate Path within an SRv6 policy.

Moreover, a segment list may be identified by more than one PSID if needed. For example, the same segment list in different Candidate Path or SR policy may use different PSIDs for identifying its path to distinguish the statistics data of other SRv6 path in other SRv6 policies using the same segment list.

This document does not define and limit the value allocation scheme for PSID, the value allocation scheme is designed depending on the use cases.

3.1. Structure of an SRv6 PSID

As mentioned above, an SRv6 PSID is stored independent from the FIB, so that the total 128-bit can be programmed in use. To avoid the value conflict, the value of a PSID MUST be global unique within the SRv6 domain.

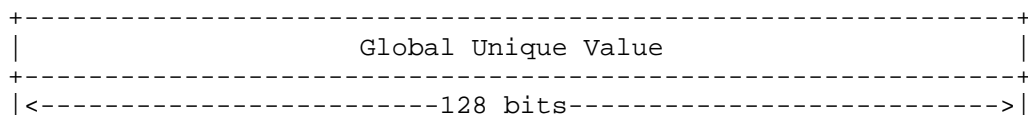


Figure 1. 128-bit PSID

A use case can define its specific structure of the PSID. In order to reuse the existing mechanism of SRv6 control plane and forwarding plane, this document defines a structure of SID reusing the structure of a SID defined in [RFC8986]. A future document MAY add further new formats for the SRv6 PSID, provided the SRv6 PSID value remains unique irrespective of the format.

3.1.1. Reusing Existing SID Structure

As per [RFC8986], an SRv6 SID consists of LOC:FUNCT:ARG, where a locator (LOC) is encoded in the L most significant bits of the SID, followed by F bits of function (FUNCT) and A bits of arguments (ARG). L, the locator length, is flexible, and an operator is free to use the locator length of their choice. F and A may be any value as long as $L+F+A \leq 128$. When $L+F+A$ is less than 128, then the remaining bits of the SID MUST be zero.

SRv6 PSID follows this structure, where the LOC part identifies the egress node that allocates the PSID, and the FUNCT part is a unique local ID to identify an SRv6 Path and its endpoint behavior, which is END.PSID (End Function with Path Segment Identifier). The code point of END.PSID is 0x0064. The Argument part is set as 0 by default. Future use cases may define the detailed usage of Arguments part.

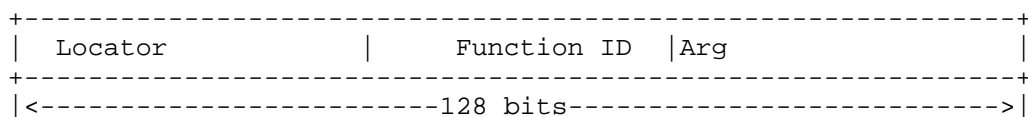


Figure 2. PSID Format following LOC:FUNCT:ARG

4. Encoding of an SRv6 PSID

This section describes the SRv6 PSID encoding in SRH.

The SRv6 PSID MUST appear only once in a segment list, and it MUST appear as the last entry in the segment list.

4.1. SRH.P-flag

To indicate the existence of a PSID in the SRH, this document defines a P-flag in the SRH flag field, and it is to be allocated, see IANA allocation section. The encapsulation of a SRv6 PSID is shown below.

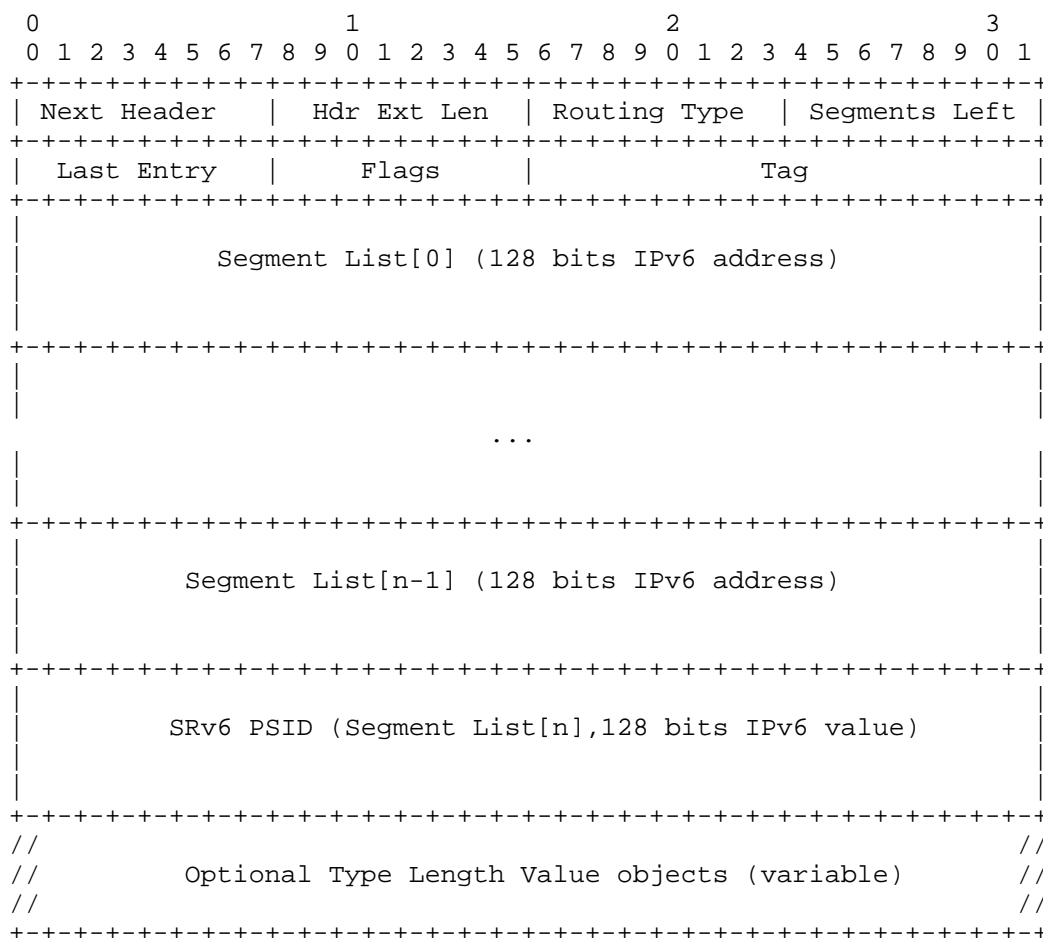


Figure 3. SRv6 PSID in SID List

- * P-flag(TBA1): set when an SRv6 PSID is inserted. A node that does not understand the P-flag will ignore it as described in [RFC8754]. A node that understands the P-flag but does not support SRv6 PSID processing MUST ignore the P-flag. If the P-flag is unset or the P-flag is ignored when processing, the SRv6 PSID processing is skipped or ignored.

SRH.P-flag processing can be enabled or disabled by configuration on devices, it can be done by CLI, NETCONF YANG or other ways, and this is out of the scope of this document.

Some use cases only require the ingress node and egress node to process the PSID, while the intermediate nodes ignore the PSID. Some use cases may require all the segment nodes along the path to process the PSID. In order to distinguish if the PSID processing is supported on the intermediate nodes, an implementation is required to support enabling and disabling the intermediate node PSID processing. The pseudo code of PSID processing is described as below. The intermediate node processing of PSID can be enabled by configuration like CLI , NETCONF YANG or other ways.

```
S01. if SRH.P-flag processing is enabled:
S02.   if intermediate node SRv6 PSID processing is disabled:
S03.     if SRH.P-flag is set and the node is the egress node: ;;ref1
S03.       SRv6 PSID processing      ;;ref2
S04.   else:
S05.     if SRH.P-flag is set:
S06.       SRv6 PSID processing
```

Ref1 : When SRv6 compression [I-D.ietf-spring-srv6-srh-compression] is not supported, SRH.SL==0 indicates the node is the egress node, which is the last segment endpoint node. When SRv6 compression is supported, an implementation identifies a node as the egress node by checking whether the SID is the last SID in the last C-SID container in the SRH as per [I-D.ietf-spring-srv6-srh-compression]. This document does not define the details of pseudo code of the implementation.

Ref2: The SRv6 PSID processing is associated with the specific application, such as SRv6 PSID based Performance measurement, so this is out of the scope of this document.

When the SRH.P-flag is set, it indicates that a PSID is encoded in the SRH and needs to be processed when processing the packet. In the cases that the intermediate processing of PSID is disabled, a node will process the PSID only when it is the last segment endpoint node. In this case, when the nodes are an intermediate node, it ignores the

processing of PSID. When the intermediate processing is enabled, all the segment endpoint nodes along the path are able to process the PSID if a PSID is encoded in the SRH. There might be some use cases that metadata of the packets need to be collected and processed on the intermediate nodes, especially for the stateful use cases. The details of these use cases are out of the scope of this document, and might be described in other documents in the future.

5. SRv6 PSID Allocation

A Path Segment is a segment on the egress node, and its identifier can be allocated through several ways, such as CLI configuration on the egress node, or allocated from the central controller by using BGP [I-D.ietf-idr-sr-policy-path-segment], PCEP [I-D.ietf-pce-sr-path-segment] or other ways. The mechanisms of allocating and distributing a PSID are out of scope of this document.

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

Depending on the use case, a PSID may be distributed to the SRv6 nodes along the SRv6 path. In this case, the SRv6 nodes that learned the PSID may process the PSID depending on the use case. This is out of the scope of this document, and may be studied in the future if needed.

6. Processing of SRv6 PSID

When the SRv6 PSID is used, the following rules also apply:

- * The SRv6 PSID MUST appear only once in a segment list, and it MUST appear as the last entry. Placing an SRv6 PSID at any other location in the SID list will result in unpredictable forwarding behavior. Only the one that appears as the last entry in the SID list will be processed.
- * The ingress needs to set the P-flag when an SRv6 PSID is inserted in the SID List. Nodes that support SRv6 PSID processing will inspect the last entry to process SRv6 PSID when the P-flag is set. When the P-flag is unset, the nodes will not inspect the last entry.

- * When an SRv6 PSID is inserted and the Last Entry is greater than 0, the SL MUST be initiated to be less than the value of Last Entry, and MUST NOT point to SRv6 PSID. For instance, when the Last entry is 4, the SID List[4] is the SRv6 PSID, so the SL MUST be set to 3 or other numbers less than Last entry.
- * When an SRv6 PSID is inserted and the PSID is the only segment in the SRH(in the case that only one NEXT-C-SID container [I-D.ietf-spring-srv6-srh-compression] is used in the IPv6 DA), then the SL and Last entry are set to 0.
- * The SRv6 PSID MUST NOT be copied to the IPv6 destination address.
- * Penultimate Segment Popping (PSP, as defined in [RFC8986]) MUST be disabled for the penultimate Segment in the SRH. In other words, a SID with PSP flavor MUST NOT be used in the penultimate segment in the SRH to avoid removing the SRH before reaching the egress node.

7. 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
0x0064	End.PSID	[This.ID]

This document also requests IANA to allocate bit position TBA1 within the "Segment Routing Header Flags" registry defined in [RFC8402].

Value	Description	Reference
TBA1	P-flag to indicate using PSID	[This.ID]

8. Security Considerations

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

Similar to SR-MPLS PSID [RFC9545], the data plane processing of a PSID is a local implementation of an SRv6 segment endpoint node, which follows the same logic of an existing SRv6 data plane.

In this document, only the egress node and the ingress node of the associated path will learn the information of a PSID. The intermediate nodes of this path will not learn it. However, in some cases, the whole Segment list with PSID may be used in a sub-set of a longer path. In this case, the whole segment list may be shared to the ingress node of the longer path. Similar to other SIDs defined in [RFC8402], the PSID must be distributed in a trusted domain under the considerations defined in Section 8.2 of [RFC8402].

A PSID is used within an SRv6 trusted domain [RFC8402] and must not leak outside the domain; therefore, no new security threats are introduced compared to current SRv6.

As per [RFC8402], SR domain boundary routers MUST filter any external traffic destined to an SID associated with a segment within the trusted domain; this applies to a PSID as well. Other security considerations of SRv6 described in Section 8.2 of [RFC8402] apply to this document. The distribution of a PSID from an egress node to an ingress node is performed within an SR-trusted domain, and it is out of the scope of this document. The details of the mechanism and related security considerations will be described in other documents.

9. 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

10. Acknowledgements

The authors would like to thank Adrian Farrel, Stefano Previdi, Zafar Ali, Lijie Deng, Zehua Hu, Joel Halpern, Yao Liu, Greg Mirsky, Quan Xiong, Changwang Lin, Weier Li and Xinyue Zhang for their valuable comments and suggestions.

11. References

11.1. 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>>.
- [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>>.
- [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>>.

11.2. Informative References

- [I-D.ietf-idr-sr-policy-path-segment] Li, C., Li, Z., Yin, Y., Cheng, W., and K. Talaulikar, "SR Policy Extensions for Path Segment and Bidirectional Path", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-path-segment-13, 2 October 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-path-segment-13>>.
- [I-D.ietf-pce-sr-bidir-path] Li, C., Chen, M., Cheng, W., Gandhi, R., and Q. Xiong, "Path Computation Element Communication Protocol (PCEP) Extensions for Associated Bidirectional Segment Routing (SR) Paths", Work in Progress, Internet-Draft, draft-ietf-pce-sr-bidir-path-15, 13 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-sr-bidir-path-15>>.

- [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)", Work
in Progress, Internet-Draft, draft-ietf-pce-sr-path-
segment-12, 14 October 2024,
<<https://datatracker.ietf.org/doc/html/draft-ietf-pce-sr-path-segment-12>>.
- [I-D.ietf-spring-srv6-srh-compression]
Cheng, W., Filsfils, C., Li, Z., Decraene, B., and F.
Clad, "Compressed SRv6 Segment List Encoding (CSID)", Work
in Progress, Internet-Draft, draft-ietf-spring-srv6-srh-
compression-23, 6 February 2025,
<<https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-srh-compression-23>>.
- [I-D.ietf-spring-stamp-srpm]
Gandhi, R., Filsfils, C., Voyer, D., Chen, M., and R. F.
Foote, "Performance Measurement Using Simple Two-Way
Active Measurement Protocol (STAMP) for Segment Routing
Networks", Work in Progress, Internet-Draft, draft-ietf-
spring-stamp-srpm-17, 8 November 2024,
<<https://datatracker.ietf.org/doc/html/draft-ietf-spring-stamp-srpm-17>>.
- [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>>.
- [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S.,
Decraene, B., Litkowski, S., and R. Shakir, "Segment
Routing with the MPLS Data Plane", RFC 8660,
DOI 10.17487/RFC8660, December 2019,
<<https://www.rfc-editor.org/info/rfc8660>>.
- [RFC9545] Cheng, W., Ed., Li, H., Li, C., Ed., Gandhi, R., and R.
Zigler, "Path Segment Identifier in MPLS-Based Segment
Routing Networks", RFC 9545, DOI 10.17487/RFC9545,
February 2024, <<https://www.rfc-editor.org/info/rfc9545>>.

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 560066
Karnataka
India
Email: dhruv.ietf@gmail.com

Yongqing Zhu
China Telecom
Guangzhou
Email: zhuyq8@chinatelecom.cn