

BFD Working Group
Internet Draft
Intended status: Informational
Expires: Oct 12, 2022

C. Lin
New H3C Technologies
W. Cheng
W. Jiang
China Mobile
April 12, 2022

S-BFD Path Consistency over SRv6
draft-lin-sbfd-path-consistency-over-srv6-01

Abstract

Bidirectional Forwarding Detection (BFD) can be used to monitor paths between nodes. Seamless BFD (S-BFD) provides a simplified mechanism which is suitable for monitoring of paths that are setup dynamically and on a large scale network. In SRv6, when a headend use S-BFD to monitor the segment list/CPath of SRv6 Policy, the forward path of control packet is indicated by segment list, the reverse path of response control packet is via the shortest path from the reflector back to the initiator (headend) as determined by routing. The forward path and reverse path of control packet are likely inconsistent going through different intermediate nodes or links. This document describes a method to keep the forward path and reverse path of S-BFD consistent when detecting SRv6 Policy.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

Internet-Draft

S-BFD Path Consistency

April 2022

This Internet-Draft will expire on October 12 2022.

Copyright Notice

Copyright (c) 2022 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 (<http://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	3
2.	Requirement for S-BFD in SRv6	3
3.	Correlate bidirectional path using Path Segment	4
4.	S-BFD Procedure with Path segment	6
4.1.	S-BFD Initiator procedure	6
4.2.	S-BFD Reflector procedure	8
5.	IANA Considerations	10
6.	Security Considerations	10
7.	References	10
7.1.	Normative References	10
	Contributors	11
	Authors' Addresses	12

[1.](#) Introduction

Segment Routing (SR) allows a headend node to steer a packet flow along any path. Per-path states of Intermediate nodes are eliminated thanks to source routing. The headend node steers a flow into an SR Policy. The packets steered into an SR Policy carry an ordered list of segments associated with that SR Policy.

S-BFD is used to monitor different kinds of paths between nodes. In SRv6, when a headend use S-BFD to monitor the segment list/CPath of SRv6 Policy, the forward and reverse path of S-BFD packet are inconsistent with high probability because the reverse path is via

IPv6 forwarding and forward path is via SRv6 segment list (loose path or explicit path).

The inconsistency impacts the detecting result. If the forward path is up and reverse path is down, then the S-BFD session will be down. If there are multiple path (segment list) in a SRv6 Policy between a headend (initiator) router and a tailend(reflector) router, multiple S-BFD session will be created for each path. Each S-BFD session uses corresponding path to send control packet, but the reverse path is identical for all S-BFD sessions. If the reverse path is down, all sessions will be down. Then the SRv6 Policy is down.

The consistency of forward and reverse path of the same S-BFD session should be guaranteed. This document describes a method to keep the forward path and reverse path of S-BFD consistent using path segment when detecting SRv6 Policy.

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.

2. Requirement for S-BFD in SRv6

Monitor SRv6 Policy using S-BFD is usually based on segment list S-BFD creates session for each segment list and associates the session with segment list.

When S-BFD initiator detects the continuity of an S-BFD session, it will use the associated segment list to encapsulate IPv6 header and SRH of the control packet.

After the reflector receives the S-BFD control packet, the response control packet should be able to return along the path to avoid the false detection of the session caused by the inconsistency of the

forward and reverse paths.

Referring to the following topology, there are two paths between Node A and D, and All nodes allocate end.x Segments. Node A and D are headend and tailend nodes of each other, and SRv6 policy is created on A and D respectively.




Figure 1: reference topology

Assuming that the deployed SRv6 policy has one candidate path and each path has two segment lists. For ease of description, segment lists with the same number on Node A and D are forward and reverse paths to each other.

Node A:

SRv6 Policy A-D
Candidate Path1
Segment list1
SID-A1, SID-B2, SID-C2
Segment list2
SID-A2, SID-E2

Node D:

SRv6 Policy D-A
Candidate Path1
Segment list1
SID-D1, SID-C1, SID-B1
Segment list2
SID-D2, SID-E1

When node A is the S-BFD initiator, S-BFD sessions for segment list1 and segment list2 could be created respectively.

The control packet of S-BFD session associated with the segment list1 is forwarded to node D according to the segment list1 of node

A. The response control packet of node D needs to be returned to node A according to the segment list1 of node D. Thus the forward and reverse paths of S-BFD packets are ensured to be consistent.

3. Correlate bidirectional path using Path Segment

A Path Segment is defined to identify an SR path in [[draft-ietf-spring-srv6-path-segment](#)]. SRv6 Path segments can be used to correlate the two unidirectional SRv6 paths at both ends of the paths.

[[draft-ietf-idr-sr-policy-path-segment](#)] proposes an extension to BGP SR Policy distribute SR policies carrying Path Segment and bidirectional path information.

Lin, et al.

Expires October, 2022

[Page 4]

Internet-Draft

S-BFD Path Consistency

April 2022

Through this extension, when distributing SRv6 policy to the headend, reverse path information and path segment of segment list can be carried together.

Node A

Node D

SRv6 Policy A-D

SRv6 Policy D-A

Candidate Path1

Candidate Path1

Segment list1

Segment list1

SID-A1, SID-B2, SID-C2

SID-D1, SID-C1, SID-B1

Path Segment: SID-Path-1

Path Segment: SID-Path-2

Reverse Path Segment:

Reverse Path Segment:

SID-Path-2

SID-Path-1

Segment list2

Segment list2

SID-A2, SID-E2

SID-D2, SID-E1

Path Segment: SID-Path-3

Path Segment: SID-Path-4

Reverse Path Segment:

Reverse Path Segment:

SID-Path-4

SID-Path-3

In this way, on the headend in both directions of the forward and reverse paths, the path segment of the paths in both directions can be obtained, and the paths in both directions use the same intermediate link.

The headend can use path segment in two directions to establish a mapping table. Using this mapping table, the headend can index the reverse path through the path segment of the forward path.

The mapping table of Node A and Node D is shown below:

Node A:

Path Segment		Reverse Path Segment	
SID-Path-1	--+	SID-Path-2	--+
SID-Path-3		SID-Path-4	--+
		segment List	
		SID-A1, SID-B2, SID-C2	--+
		SID-A2, SID-E2	--+

Node D:

Path Segment		Reverse Path Segment	
SID-Path-2	--+	SID-Path-1	--+
SID-Path-4		SID-Path-3	--+
		segment List	
		SID-D1, SID-C1, SID-B1	--+
		SID-D2, SID-E1	--+

Figure 2: mapping table

4. S-BFD Procedure with Path segment

This document proposes to forward S-BFD control packets and response control packets through the consistent path by path segment.

[4.1.](#) S-BFD Initiator procedure

For instance, the S-BFD initiator is Node A in Figure 1, and the S-BFD session is bounded with Segment List1 of Policy A-D. The encapsulation format of S-BFD control packet is as follows:

```
+-----+
| IPv6 Header                                     |
. Source IP Address = S-BFD Initiator IPv6 Address .
. Destination IP Address = SegmentList[SL]         .
. Next-Header = SRH (43)                           .
.                                                    .
+-----+
| SRH as specified in RFC 8754                     |
. Next-Header = IPv6                               .
. <PathSegment, Segment List>                      .
.                                                    .
+-----+
| IPv6 Header                                     |
. Source IP Address = S-BFD Initiator IPv6 Address .
. Destination IP Address = S-BFD Reflector IPv6 Address .
. Next-Header = UDP                                 .
.                                                    .
+-----+
| UDP Header                                     |
```




Figure 3: Encapsulation format of S-BFD control packet

NodeA Encapsulates the path segment of segment list1 in SRH, and set SRH.P-Flag.

The S-BFD control packet is as follows:

A----->B----->C----->D

```
+-----+
| SA=A's Ipv6Addr |
+-----+
| DA=SID-A1       |
+-----+
| SL=3 | P-Flag=1 |
+-----+
| D's ipv6Addr    |
+-----+
| SID-C2          |
+-----+
| SID-B2          |
+-----+
```

```
+-----+
| SA=A's Ipv6Addr |
+-----+
| DA=D's ipv6Addr |
+-----+
| SL=0 | P-Flag=1 |
+-----+
| D's ipv6Addr    |
+-----+
| SID-C2          |
+-----+
| SID-B2          |
+-----+
```





Figure 4: Example of S-BFD control packet

4.2. S-BFD Reflector procedure

S-BFD control packet is forwarded along the path A->B->C->D. While packet arrives at Node D, RH.SL is 0 and the destination address is IPv6 address of Node D. Packet is delivered up to the S-BFD module in control plane.

S-BFD module detects SRH.P-flag is set, extracts the path segment of the forward path from SRH, gets the path segment of the reverse path through the mapping table. When responding to S-BFD control packet, S-BFD module uses the segment list associated with path segment of the reverse path to encapsulate SRH.

The encapsulation format of S-BFD response control packet is as follows:

```

+-----+
| IPv6 Header |
. Source IP Address = S-BFD Reflector IPv6 Address .
. Destination IP Address = SegmentList[SL] .
. Next-Header = SRH (43) .
. .
+-----+
| SRH as specified in RFC 8754 |
. Next-Header = IPv6 .
. <Segment List> .

```

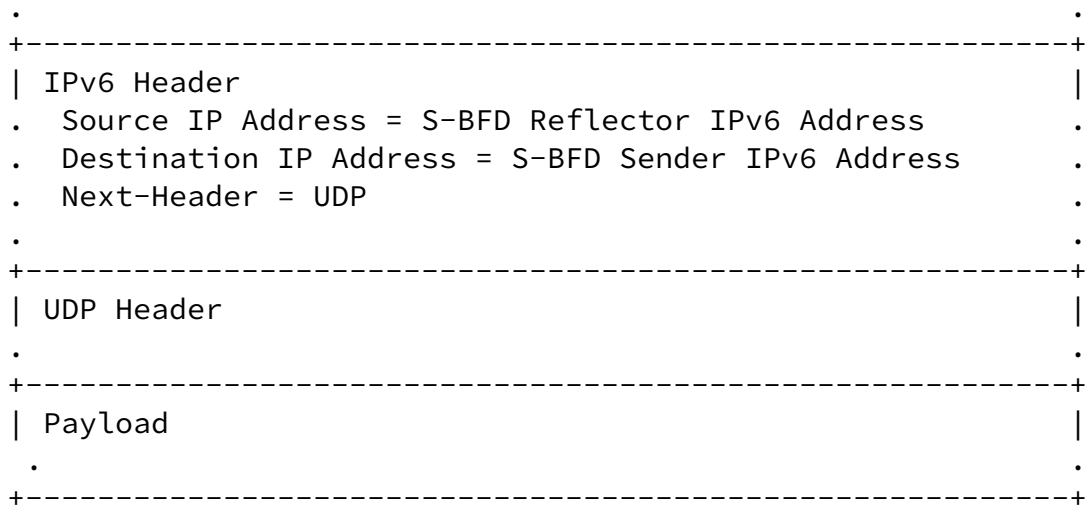



Figure 5: Encapsulation format of S-BFD response control packet

The Example of S-BFD response control packet is as follows:



SL=3 P-Flag=0	SL=0 P-Flag=0
+-----+	+-----+
A's ipv6Addr	A's ipv6Addr
+-----+	+-----+
SID-B1	SID-B1
+-----+	+-----+
SID-C1	SID-C1
+-----+	+-----+
SID-D1	SID-D1
+-----+	+-----+
sbfd-payload	sbfd-payload
+-----+	+-----+

Figure 6: Example of S-BFD response control packet

The S-BFD response control packet will be forward along the path D->C->B->A. In this way, the forward and reverse paths of S-BFD are guaranteed to be consistent.

5. IANA Considerations

This document has no IANA actions.

6. Security Considerations

The security requirements and mechanisms described in [[RFC8402](#)] and [[RFC8754](#)] also apply to this document.

This document does not introduce any new security consideration.

7. References

7.1. Normative References

[I-D.ietf-idr-segment-routing-te-policy] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Rosen, E., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", [draft-ietf-idr-segment-routing-te-policy-14](#) (work in progress), November 2021

Segment Routing Network", [draft-ietf-spring-mpls-path-segment-07](#) (work in progress), December 2021.

[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-18](#) (work in progress), February 2022.

[I-D.ietf-spring-srv6-path-segment] Li, C., Cheng, W., Chen, M., Dhody, D., and Y. Zhu, "Path Segment for SRv6 (Segment Routing in IPv6)", [draft-ietf-spring-srv6-path-segment-03](#) (work in progress), November 2021.

[I-D.ietf-idr-sr-policy-path-segment] Li, C., Li, Z., Yin, Y., Cheng, W., Talaulikar, K., "SR Policy Extensions for Path Segment and Bidirectional Path", [draft-ietf-idr-sr-policy-path-segment-05](#) (work in progress), January 2022.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", [RFC 5880](#), DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.

[RFC7880] Pignataro, C., Ward, D., Akiya, N., Bhatia, M., and S. Pallagatti, "Seamless Bidirectional Forwarding Detection (S-BFD)", [RFC 7880](#), DOI 10.17487/RFC7880, July 2016, <<https://www.rfc-editor.org/info/rfc7880>>.

[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>>.

Contributors

Yisong Liu contributed to the content of this document.

Authors' Addresses

Changwang Lin
New H3C Technologies
Beijing
China

Email: linchangwang.04414@h3c.com

Weiqiang Cheng
China Mobile
Beijing
CN

Email: chengweiqiang@chinamobile.com

Wenying Jiang
China Mobile
Beijing
CN

Email: jiangwenying@chinamobile.com

