

SPRING Working Group
Internet Draft
Intended status: Standards Track
Expires: February 05, 2026

L. Gong
China Mobile
C. Lin
Y. Qiu
New H3C Technologies
X. Min
ZTE Corporation
August 04, 2025

Ping Path Consistency over SRv6
draft-gong-spring-ping-path-consistency-over-srv6-05

Abstract

In the SRv6 network, the headend node could use Ping (ICMPv6 Echo) to detect the connectivity of the SRv6 path to implement path switching. When a headend use Ping to detect the segment list/CPath of SRv6 Policy, the forward path of ICMPv6 Echo Request message is indicated by segment list, the reverse path of ICMPv6 Echo Reply message is via the shortest path from the destination node to the source as determined by routing. The forward path and reverse path of ICMPv6 message are likely inconsistent going through different intermediate nodes or links. This document describes how to ensure the consistency of the forward path and the reverse path when using ICMPv6 Echo messages to detect 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>

This Internet-Draft will expire on February 05 2026.

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 (<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	3
1.1. Requirements Language	3
2. Requirement of Path Consistency for Ping	4
3. Correlate Bidirectional Path Using Path Segment	5
3.1. procedure of Ping Detection Initiator	7
3.2. Procedure of Ping's destination Node	8
4. IANA Considerations	10
5. Security Considerations	10
6. References	10
6.1. Normative References	10
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.

SR can be instantiated on the MPLS data plane (MPLS-SR) and the IPv6 data plane (SRv6). On the SRv6 data plane, a segment is encoded as an IPv6 address (SRv6 SID) [RFC8986], and an ordered list of segments is encoded as an ordered list of SRv6 SIDs in the SR header (SRH) [RFC8754].

In the SRv6 network, the headend node could use Ping (ICMPv6 Echo) to detect the connectivity of the SRv6 path to implement path switching. When a headend use Ping to detect the segment list/CPath of SRv6 Policy, the forward path of ICMPv6 Echo Request message is indicated by segment list, the reverse path of ICMPv6 Echo Reply message is via the shortest path from the destination node to the source as determined by routing. The forward path and reverse path of ICMPv6 message are likely inconsistent going through different intermediate nodes or links.

The inconsistency impacts the detecting result. If the forward path is up and reverse path is down, then the ICMPV6 Echo Reply message will not be able to be sent to the source, and Ping detection will timeout. If there are multiple path (segment list) in a SRv6 policy between a headend (local system) node and a tailend (remote system) node, Ping detection instance will be created for each path. Each Ping detection instance uses corresponding path to send ICMPV6 Echo Request message, but the reverse path is identical. If the reverse path is down, all Ping detection instances will be down. Then the SRv6 policy is down.

The transmission path of forward ICMPV6 Echo Request messages and reverse ICMPV6 Echo Reply messages for the same Ping detection instance should be consistent.

This document describes how to ensure the consistency of the forward path and the reverse path when using ICMPV6 Echo messages to detect 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 of Path Consistency for Ping

Detecting the connectivity of SRv6 policy using Ping is usually based on segment list. Ping detection instance is created for each segment list and associated with the segment list.

Referring to the topology in Figure 1, there are two paths between Node A and D, and All nodes allocate end.x Segments on SRv6 data plane or adjacency SIDs on SR-MPLS data plane. Node A and D are headend and tailend nodes of each other, and SRv6 policy is created on A and D respectively.

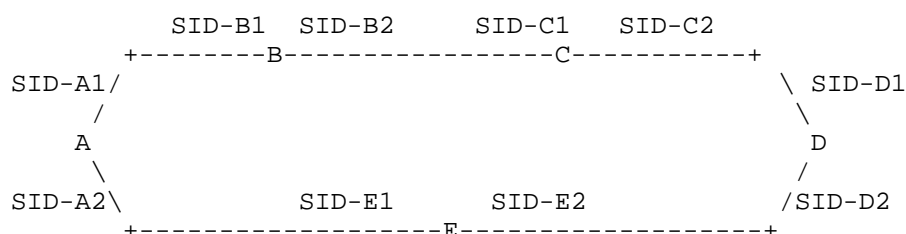


Figure 1

Assuming that the deployed SRv6 policy has one candidate path and each path has two segment lists.

Node A:	Node D:
SR Policy A-D	SR Policy D-A
Candidate Path1	Candidate Path2
Segment list1-1	Segment list2-1
SID-A1, SID-B2, SID-C2	SID-D1, SID-C1, SID-B1
Segment list1-2	Segment list2-1
SID-A2, SID-E2	SID-D2, SID-E1

List1-1 and List2-1 correspond to a forward and reverse SR forwarding path between node A and node D. List2-1 and List2-2 also correspond to another forward and reverse forwarding path between node A and node D.

Both Node A and Node D serve as head nodes and need to detect the connectivity of the segment list of a SRv6 policy.

When node A is the Ping detection initiator, Ping detection instances for segment list1 and segment list2 will be created respectively. Node A uses the associated segment list to encapsulate IPv6 header and SRH of the ICMPV6 Echo Request messages.

After Node D receives an ICMPV6 Echo Request message to its local address (or SID), the ICMPV6 Echo Reply message should be able to return along the same path to avoid the false detection of the Ping instance caused by the inconsistency of the forward and reverse paths.

The ICMPV6 Echo Request messages associated with the segment list1-1 is forwarded to node D according to the segment list1-1 of node A. The corresponding ICMPV6 Echo Reply messages of node D need to be sent to node A according to the segment list2-1 of node D. Thus, the forward and reverse paths of ICMPV6 message are ensured to be consistent.

3. Correlate Bidirectional Path Using Path Segment

This draft proposes a method to achieve consistency in the forward and reverse paths of Ping detection packets using path segment.

A Path Segment is defined to identify an SR path. In SR for MPLS data plane (SR-MPLS), Path Segment is defined in [I-D.ietf-spring-mpls-path-segment]. In SR for IPv6 data plane (SRv6), Path Segment is defined in [I-D.ietf-spring-srv6-path-segment].

Path segments can be used to correlate the two unidirectional SR paths at both ends of the paths.

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

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

Node A	Node D
SR Policy A-D	SR Policy D-A
Candidate Path1	Candidate Path2
Segment list1-1	Segment list2-1
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 list1-2	Segment list2-2
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 node 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 links.

The headend node can use path segment in two directions to establish a mapping table. Using this mapping table, the headend node can get 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:

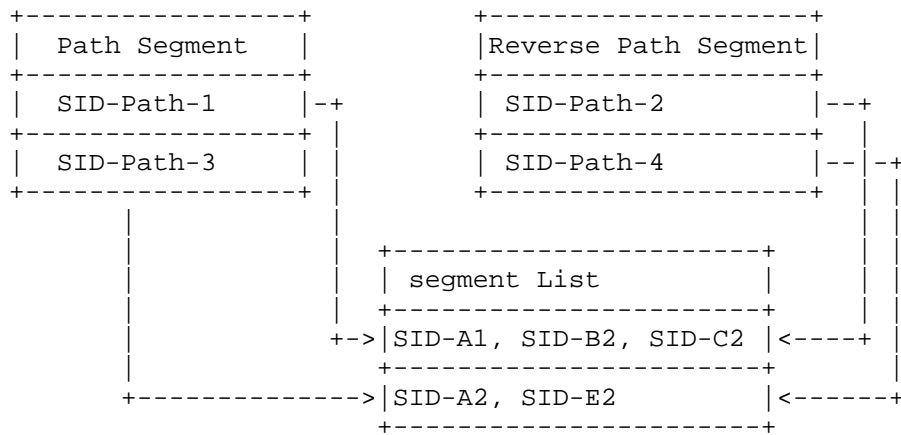


Figure 2: Mapping Table on Node A

Node D:

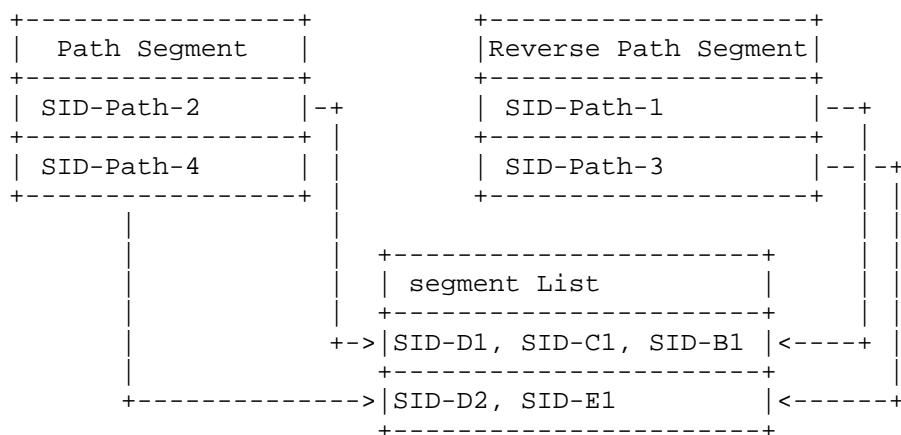


Figure 3: Mapping Table on Node D

For instance, the Ping detection initiator is Node A in Figure 1, and the Ping detection instance is bounded with Segment List1-1 of Policy A-D.

3.1. procedure of Ping Detection Initiator

When path segment is used, the encapsulation format of ICMPV6 Echo Request message from Node A to Node D is as follows:

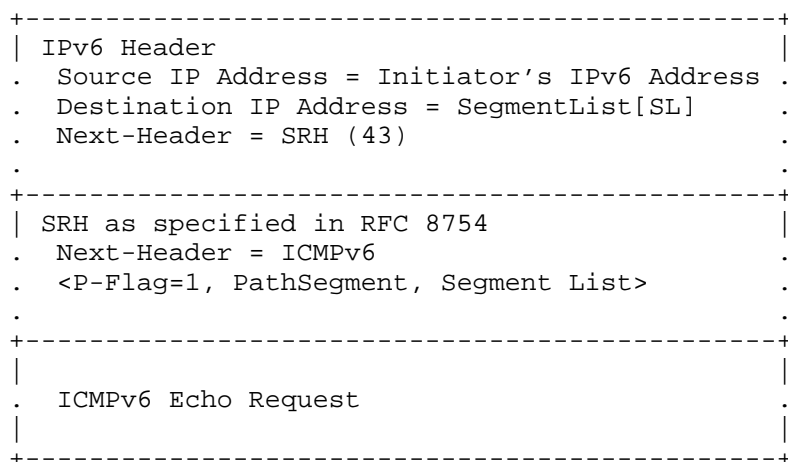


Figure 4: ICMPv6 Echo Request Message Format

Node A Encapsulates the path segment of segment list1-1 in SRH and sets SRH.P-Flag.

The ICMPv6 Echo Request message is encapsulated and forwarded as follows:

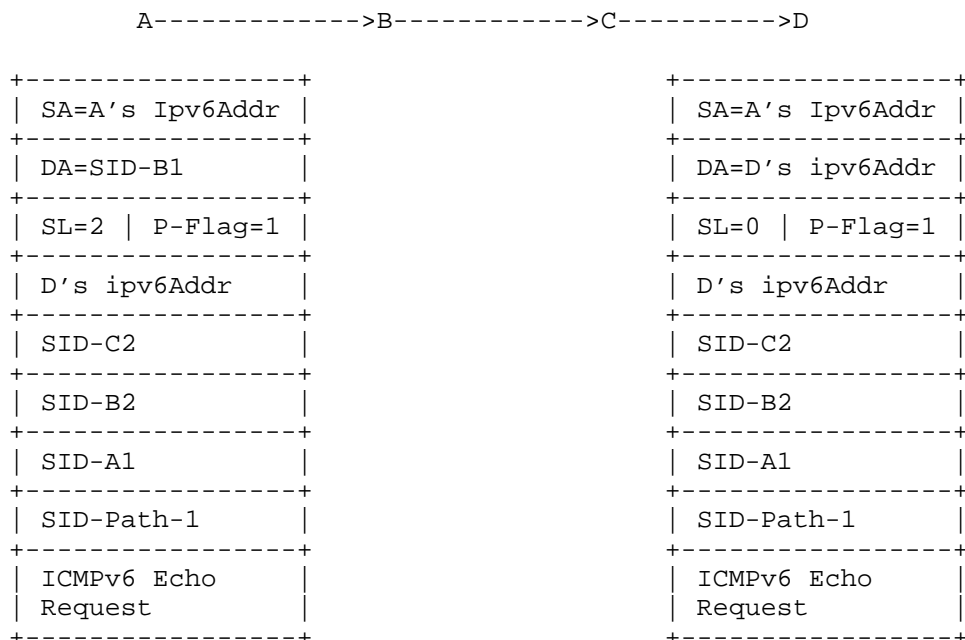


Figure 5: Example of ICMPv6 Echo Request Message

3.2. Procedure of Ping's destination Node

ICMPv6 Echo Request message is forwarded along the path A->B->C-D. While packet arrives at Node D, SRH.SL is 0 and the destination address is IPv6 address of Node D. Packet is delivered up to control plane.

Control plane 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. Then use the segment list associated with path segment of the reverse path to encapsulate SRH of ICMPv6 Echo Reply message.

The encapsulation format of ICMPv6 Echo Reply message is as follows:

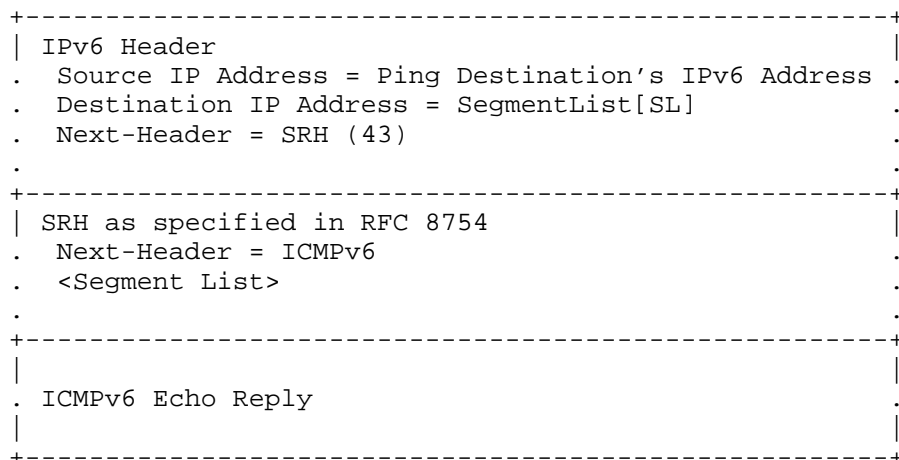


Figure 6: ICMPv6 Echo Reply Message Format

The ICMPv6 Echo Reply message is encapsulated and forwarded as follows:

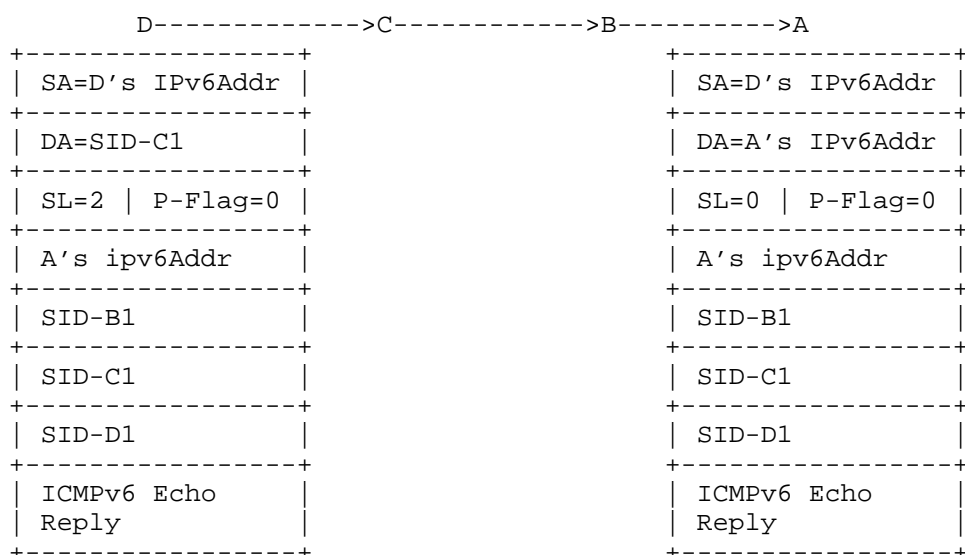


Figure 7: Example of ICMPv6 Echo Reply Message

The ICMPv6 Echo Reply message will be forward along the path D->C->B->A. In this way, the forward and reverse paths of ICMPV6 packets are guaranteed to be consistent.

4. IANA Considerations

This document has no IANA actions.

5. 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.

6. References

6.1. Normative References

- [I-D.ietf-idr-segment-routing-te-policy] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", draft-ietf-idr-segment-routing-te-policy-20 (work in progress), July 2022
- [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-08 (work in progress), September 2022.
- [I-D.ietf-spring-srv6-path-segment] Li, C., Cheng, W., Chen, M., Dhody, D., and Zhu, Y., "Path Segment for SRv6 (Segment Routing in IPv6)", draft-ietf-spring-srv6-path-segment-06 (work in progress), May 2023.
- [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-07(work in progress), February 2023.
- [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>>.

Authors' Addresses

Liyan Gong
China Mobile
China

Email: gongliyan@chinamobile.com

Changwang Lin
New H3C Technologies
China

Email: linchangwang.04414@h3c.com

Yuanxiang Qiu
New H3C Technologies
China

Email: qiuyuanxiang@h3c.com

Xiao Min
ZTE Corporation
Nanjing
China

Email: xiao.min2@zte.com.cn

