SRv6OPS
Internet Draft
Intended status: Standards Track
Expires: April 28, 2026

L. Gong, Ed.
China Mobile
C. Lin, Ed.
New H3C Technologies

October 20, 2025

SRv6 OAM Deployment Consideration
draft-gong-srv6ops-srv6-oam-deployment-01

Abstract

This document introduces common issues to consider when implementing SRv6 OAM, as well as various solutions.

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 April 28, 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 (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....3
2. Terminology....3
3. Deployment Options....3
3.1. ping/tracert SRv6 sid....3
3.2. ping/tracert SRv6 locator....3
3.3. ping/tracert segment-list....4
3.4. ping/tracert SRv6 Policy....4
4. Considerations....4
4.1. SRv6 BE/TE....4
4.2. Encap Mode....5
4.3. EVPN....6
4.4. Locations....6
4.5. Path MTU....7
4.6. Inconsistency paths....7
4.7. Network Slicing....7
4.8. BSID stitching....8
5. IANA Considerations....8
6. Security Considerations....8
7. Acknowledgements....8
8. References....8
8.1. Normative References....8
8.2. Informative References....9
Contributors....9
Authors' Addresses....10

1. Introduction

Segment Routing IPv6 (SRv6) [RFC8986] is a network architecture that leverages IPv6 data plane encapsulation to enable flexible and efficient traffic engineering.

[I-D.liu-srv6ops-problem-summary] provides an overview of the common problems encountered during SRv6 deployment and operation. It provides a foundation for further work, including potential solutions and best practices to navigate deployment.

SRv6 OAM (Operations, Administration, and Maintenance) is used to detect the connectivity of SRv6 paths and locate faults in SRv6 paths. SRv6 Ping is used to check network connectivity and host reachability. SRv6 Tracert can verify network connectivity and also analyze where a network fault has occurred.

This document introduces common issues to consider when implementing SRv6 OAM, as well as various solutions.

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. Terminology

SRv6: Segment Routing based on IPv6

OAM: Operations, Administration, and Maintenance

NSH: Network Service Header

3. Deployment Options

This chapter describes several methods that must be supported when implementing SRv6 OAM.

3.1. ping/tracert SRv6 sid

SRv6 Ping verifies reachability by sending ICMP requests to the target SID and waiting for replies, while SRv6 Traceroute uses incrementing Hop Limit values to probe the path hop-by-hop, constructing the path topology and locating faults by analyzing timeout or destination unreachable messages from each hop. Both rely on the basic IPv6 forwarding mechanism, but Traceroute additionally provides path visibility, making it suitable for network outage troubleshooting.

3.2. ping/tracert SRv6 locator

SRv6 Ping/Traceroute for Locator operates similarly to SID-based verification but targets an entire SRv6 locator prefix instead of individual SIDs. The source node sends ICMP probes (Ping) or Hop Limit-graduated packets (Traceroute) destined to any address within the locator’s range, allowing network operators to validate routing consistency across all nodes sharing that locator block. While Ping confirms basic reachability of the locator space, Traceroute reveals the actual forwarding path through the SRv6 domain, helping identify potential routing asymmetries or black holes within the advertised prefix scope.

3.3. ping/tracert segment-list

SRv6 Segment List Ping/Traceroute verifies the reachability and forwarding status of explicit paths by constructing probe packets containing specified segment lists: the Ping operation sends ICMP request messages with a complete Segment Routing Header (SRH), requiring the end point node to return a reply to confirm the validity of the entire path; Traceroute, on the other hand, uses a progressive probing method, successively truncating the segment list and leveraging hop-by-hop returned ICMP Time Exceeded messages to reconstruct the actual forwarding path, thereby verifying service link integrity and precisely identifying faulty nodes in multi-segment routes. This mechanism is particularly suitable for verifying complex cross-domain SRv6 policies and explicit paths issued by SDN controllers.

3.4. ping/tracert SRv6 Policy

SRv6 Policy Ping/Traceroute validates the implementation of SRv6 policy paths by generating test packets that follow predefined segment lists. For Ping, the source node constructs an IPv6 packet with an SRH containing the full policy path segments, where successful end-to-end ICMP reply confirms policy execution correctness. Traceroute progressively tests each segment by incrementally expanding the active segment list length while monitoring intermediate node responses through ICMP time exceeded messages, enabling both path visualization and fault isolation within the policy-defined route. This approach effectively verifies Traffic Engineering policies and detects misconfigurations in SRv6-enabled networks.

4. Considerations

This chapter describes the key scenarios that SRv6 OAM must support to ensure reliable network operation and accurate fault detection.

4.1. SRv6 BE/TE

When implementing SRv6 OAM, simultaneous support for SRv6 BE and SRv6 TE must be considered.

In cases where an SRv6 Policy contains multiple Segment-List paths, Ping/Traceroute tests should be individually initiated for each Segment-List, for example by assigning different Flow Labels or DSCP markings to distinguish the paths. When testing multiple paths, the system should preferably support parallel testing to improve efficiency.

By using PING BE/TE, the SRv6 OAM detection method can be divided into segment-by-segment detection and non-segment-by-segment detection based on path coverage. Segment-by-segment detection refers to verifying connectivity between the source node and all SRv6 nodes along the SRv6 forwarding path. Non-segment-by-segment detection refers to only verifying connectivity between the source node and the destination node. There are significant differences in design objectives and implementation logic between the two.

Segment-by-segment detection enables fine-grained fault localization and is suitable for troubleshooting intermittent packet loss segment by segment. Non-segment-by-segment detection only verifies end-to-end connectivity between the source and destination nodes, without concern for the state of intermediate nodes.

4.2. Encap Mode

When implementing SRv6 OAM (Operations, Administration, and Maintenance), it is essential to consider various testing scenarios, including Ping/Traceroute verification for both SRv6 basic forwarding paths (BE, Best Effort) and SRv6 TE Policies (Traffic Engineering Policy). The choice of encapsulation mode significantly impacts detection accuracy and efficiency, with the following key considerations:

Encap (Encapsulation) Mode: Suitable for end-to-end path validation, where a new IPv6 header and outer SRH (Segment Routing Header) are encapsulated around the original packet. This ensures intermediate nodes process the packet according to the specified SID list, making it ideal for verifying explicit path correctness.

Insert (Insertion) Mode: Instead of adding a new IPv6 header, this mode directly inserts an SRH into the existing IPv6 header. It is typically used in dynamic path adjustment scenarios, such as inserting additional SIDs at specific nodes in Service Chaining.

Compressed Mode (e.g., uSID): To reduce SRH overhead and improve transmission efficiency, compressed SID formats (e.g., micro-SID/uSID) can be adopted. This is particularly critical in low-bandwidth or high-throughput environments like 5G transport networks or data center interconnects.

Furthermore, practical deployment must account for compatibility across encapsulation modes and select appropriate probing strategies based on network conditions to ensure optimal OAM performance.

4.3. EVPN

A solution for implementing EVPN Ping over SRv6 needs to be considered.

When executing EVPN Ping in an SRv6 network, the encapsulation format should be selected according to path characteristics. Currently, two feasible frameworks exist: Nested SRv6 (L3-in-L3), with the format [Ethernet] [IPv6 (Outer)] [SRH] [IPv6 (Inner)] [UDP] [MPLS Ping Payload], and L2-over-SRv6 (Ethernet over SRv6), with the format [Ethernet (Outer)] [IPv6] [SRH] [Ethernet (Inner)] [IPv6] [UDP] [MPLS Ping].

4.4. Locations

In SRv6 networks, Ping/Traceroute implementations vary significantly depending on the network role and visibility scope of the measurement point (MP). When implementing SRv6 OAM, support must be ensured for initiating tests from all locations.

Below are the layered detection requirements in a typical CE-PE-P-P-PE-CE scenario:

Detection initiated from CE side

Can only verify end-to-end service reachability

Encapsulation requirements: standard IPv6 Ping (without SRH) or End.DT4/SID with target PE

Limitation: unable to perceive internal SRv6 path details within the carrier network

Detection initiated from PE devices

Supports two modes:

a) Service-layer detection: simulates CE perspective using End.DX4/SID encapsulation

b) Tunnel-layer detection: verifies TE path with complete SRH list

Key capability: must support VRF-aware Ping (to distinguish between tenants)

Detection from P devices (intermediate nodes)

Core requirements:

Must support SRH parsing and processing

Must enable ICMPv6 error message generation

Special scenarios:

Decoding capability for compressed SIDs (uSID)

Support trace validation for TI-LFA protected paths

Special handling for border nodes

Between PE and P, support is required for:

Explicit path verification (with specified Segment List)

Implicit path verification (relying on IGP computation)

Fault localization: determine the breakpoint using the "Segments Left" field in the SRH

4.5. Path MTU

When implementing SRv6 OAM, support for dynamically discovering Path MTU via SRv6 OAM should be considered. The source node initiates the probe and collects the effective MTU of the path.

4.6. Inconsistency paths

During the implementation of SRv6 OAM (Operations, Administration, and Maintenance), the inconsistency between forward and reverse paths caused by asymmetric routing is one of the key challenges affecting network reliability and diagnostic accuracy. This can lead to inaccurate OAM measurements and difficulty in fault localization. This issue needs to be carefully addressed during implementation.

4.7. Network Slicing

When implementing SRv6 OAM, it is necessary to support the network slicing function to ensure that different service slices can obtain independent operations and maintenance support. When a slice SID path fails, it should trigger a fast switchover of the corresponding slice path without affecting other slices.

4.8. BSID stitching

When performing SRv6 TE Policy Tracert operations, if the SID list of an SRv6 TE Policy contains another SRv6 TE Policy's BSID (for example, when SRv6 TE Policy A's SID list stitches together with SRv6 TE Policy B's BSID), this scenario is referred to as BSID stitching. Special considerations must be made for such implementations.

In typical scenarios, since nodes in SRv6 TE Policy B's SID list will only send ICMPv6 Time Exceeded messages back to SRv6 TE Policy B's headend node rather than SRv6 TE Policy A's headend node, this may cause SRv6 TE Policy A's headend to incorrectly detect unreachable nodes within its own policy.

To address this issue, the ICMPv6 error message relay function can be enabled on SRv6 TE Policy B's headend node. When receiving ICMPv6 Time Exceeded messages from other nodes in SRv6 TE Policy B, the headend node of Policy B will act as a proxy node to forward these ICMPv6 Time Exceeded messages to the headend node of SRv6 TE Policy A, thereby preventing false detection.

5. IANA Considerations

TBD.

6. Security Considerations

TBD.

7. Acknowledgements

TBD.

8. References

8.1. Normative References

[RFC7307] Q. Zhao, Huawei Technology, K. Raza, C. Zhou, Cisco Systems, L. Fang, Microsoft, L. Li, China Mobile, D. King, Old Dog Consulting, "LDP Extensions for Multi-Topology", RFC 5286,DOI 10.17487/RFC7307, July 2014,<http://www.rfc-editor.org/info/rfc7307>.

[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>.

[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>.

8.2. Informative References

[I-D.liu-srv6ops-problem-summary]Liu, Y., Graf, T., Miklos, Z., Contreras, L. M., and N. Leymann, "SRv6 Deployment and Operation Problem Summary", Work in Progress, Internet-Draft, draft-liu-srv6ops-problem-summary-06, 26 September 2025,<https://datatracker.ietf.org/doc/html/draft-liu-srv6ops-problem-summary-06>.

Contributors

Authors' Addresses

Liyan Gong (editor)
China Mobile
China
Email: gongliyan@chinamobile.com

Changwang Lin (editor)
New H3C Technologies
China
Email: linchangwang.04414@h3c.com