

SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: 1 November 2025

L. Han
China Mobile
J. Dong
Huawei Technologies
M. Wang
China Mobile
R. Chen
ZTE Corporation
Z. Du
China Mobile
30 April 2025

SRv6 for Inter-Layer Network Programming
draft-ietf-spring-srv6-inter-layer-programming-00

Abstract

The Segment Routing over IPv6 (SRv6) Network Programming framework enables a network operator or an application to specify a packet processing program by encoding a sequence of instructions in the IPv6 packet header.

Following the SRv6 Network Programming concept, this document defines SRv6 based mechanisms for inter-layer network programming, which can help to integrate the packet network layer with its underlying layers efficiently. For inter-layer path programming, a new SRv6 behavior is defined for steering packets to underlay network connections. The applicability of this new SRv6 behavior in typical scenarios is illustrated.

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 1 November 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	3
2. Use Cases of Inter-Layer Network Programming	3
2.1. IP and Optical Inter-layer Path Programming	3
2.2. IP and MTN Inter-layer Path Programming	3
3. SRv6 Behavior for Inter-Layer Path Programming	4
4. Application of SRv6 Inter-Layer Programming	6
4.1. IP and Optical Integration	6
4.2. IP and MTN Integration	8
5. Security Considerations	10
6. IANA Considerations	10
7. Acknowledgements	11
8. References	11
8.1. Normative References	11
8.2. Informative References	11
Authors' Addresses	12

1. Introduction

In many network scenarios, the operator owns a multi-layered network. The technology in layer-3 has converged to IP, while there can be different technologies in layer-2 and below. In such networks, the cross-layer planning and optimization is considered more efficient than independent planning and operation of the layer-3 and the underlying networks in terms of resource utilization and SLA assurance, but it is also considered more complicated. Thus a mechanism for flexible and efficient inter-layer network integration is desired.

Segment Routing over IPv6 (SRv6) [RFC8986] enables a network operator or an application to specify a packet processing program by encoding a sequence of instructions in the IPv6 packet header. Currently SRv6

does not consider about the network layers under the IP layer. However, with the capability of SRv6 network programming, it is possible to achieve seamless integration between IP (layer-3) and the underlying (layer-2 and below) networks.

Following the SRv6 network programming concept, this document defines a new SRv6 behavior, which can be used for steering packets to underlay network connections, so that the packet network layer can be integrated with the underlying layers efficiently. The applicability of this new SRv6 behavior in typical inter-layer network programming scenarios is also illustrated. The proposed mechanism is applicable to networks in which the IP layer and its underlay network are under the same administration, and the necessary information of the underlay connections is allowed to be exposed to the IP layer for inter-layer path programming.

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. Use Cases of Inter-Layer Network Programming

2.1. IP and Optical Inter-layer Path Programming

In many network scenarios, the underlay of the IP network is an optical network. The IP network and optical network are usually managed separately, the optical network works as an underlay which is normally invisible to the IP network. In some cases, the IP layer path and the underlay optical resources may not be one-to-one mapping, which makes the redundant optical paths not fully used by the IP layer. In some other cases, there may be optical paths between non-adjacent IP nodes, thus they are not visible in the L3 topology and can not be used for carrying traffic based on IP routing. However, such optical paths may be used for inter-layer traffic engineering.

2.2. IP and MTN Inter-layer Path Programming

The architecture of Metro Transport Network (MTN) is defined in [ITU-T_G.8310]. In an MTN based network, network nodes can support two forwarding modes: per-hop IP packet forwarding and the MTN Path (MTNP) layer cross-connect. An MTN path is a multi-hop underlay transport path which may be established between any two nodes in the MTN network, and the intermediate nodes on the MTN path will forward

the traffic based on the pre-established MTN cross-connect without IP table lookup. Thus an MTN path is considered as an underlay connection between two remote MTN nodes. Although in some cases it is possible to set up a layer-3 adjacency between the two endpoints of the MTN path, it will make the provisioning of MTN path complicated. Moreover, in some cases the two endpoints may reside in different IGP areas or ASes, which makes a layer-3 adjacency between them more challenging. Last but not the least, an MTN path may be provisioned unidirectionally, which cannot pass the bidirectional connectivity check required for a layer-3 link. Since the MTN paths are usually not visible in the L3 topology, it is difficult to compute and establish an end-to-end inter-layer path which consists of both the layer-3 network segments and the MTN paths.

3. SRv6 Behavior for Inter-Layer Path Programming

In this section, a new SRv6 Endpoint Behavior is proposed for SRv6 based inter-layer path programming.

The "Endpoint with inter-layer connection" behavior ("End.IL" for short) is a variant of the SRv6 End behavior as defined in [RFC8986]. Its main use is for SRv6 based inter-layer path programming and traffic engineering. The End.IL behavior steers packets to a remote network node via underlay network connections.

The underlay network connections may be realized using Metro Transport Network (MTN) paths [ITU-T_G.8310], ODUK or DWDM connections, or other technologies which work as the underlay of the IP network. Usually the underlay connections are invisible in the L3 topology and are not considered in IP based distributed route computation (e.g. SPF). However, this is just the expected behavior in some inter-layer programming scenarios, where the underlay connections are provisioned for traffic engineering of specific types of services. The SRv6 End.IL SIDs can be used together with other types of SRv6 SIDs to build SRv6 SID lists for inter-layer path programming.

Any SID instance of this behavior is associated with an underlay network connection, which connects to a remote network node.

When node N receives a packet destined to S and S is a local End.IL SID, N does the following:

```
S01. When an SRH is processed {
S02.   If (Segments Left == 0) {
S03.     Stop processing the SRH, and proceed to process the next
        header in the packet, whose type is identified by
        the Next Header field in the routing header.
S04.   }
S05.   If (IPv6 Hop Limit <= 1) {
S06.     Send an ICMP Time Exceeded message to the Source Address
        with Code 0 (Hop limit exceeded in transit),
        interrupt packet processing, and discard the packet.
S07.   }
S08.   max_LE = (Hdr Ext Len / 2) - 1
S09.   If ((Last Entry > max_LE) or (Segments Left > Last Entry+1)) {
S10.     Send an ICMP Parameter Problem to the Source Address
        with Code 0 (Erroneous header field encountered)
        and Pointer set to the Segments Left field,
        interrupt packet processing, and discard the packet.
S11.   }
S12.   Decrement IPv6 Hop Limit by 1
S13.   Decrement Segments Left by 1
S14.   Update IPv6 DA with Segment List[Segments Left]
S15.   Send the packet through the underlay network connection
        identified by S.
S16. }
```

Note that the underlay network connection in step 15 SHOULD be established before the associated End.IL SID is announced into the network.

When forwarding packets through the underlay network connection towards the remote endpoint node, depending on the type of connection, layer-2 encapsulation may be required. The information needed for layer-2 encapsulation may be provisioned via mechanisms such as static Neighbor Discovery (ND) cache or using specific MAC address. In some network scenarios where there is only one underlay network connection between two layer-3 nodes, the use of End.IL SID could be optional. In some implementations, the End.IL SID may optionally contain the arguments field, which can be used to encode the type and identifier information of the underlay connection. The details are out of the scope of this document.

End.IL SIDs MAY be allocated by the network nodes and announced to the network controller with the information of the underlay connections using IGP [RFC1195] [RFC2328] or BGP-LS [RFC9552]. Alternatively, End.IL SIDs MAY be assigned by a centralized network controller and advertised to the network nodes through Netconf/YANG

[RFC6241] [RFC6020], BGP [RFC4271], or PCEP [RFC5440]. The detailed protocol extensions are out of the scope of this document, and will be described in separate documents. With the information of End.IL SIDs, the network controller or headend nodes could use End.IL SIDs together with other types of SRv6 SIDs to build SRv6 SID lists for inter-layer TE path programming.

4. Application of SRv6 Inter-Layer Programming

4.1. IP and Optical Integration

Assuming that an operator owns both the IP and optical network, and the operator needs to deploy E2E service across IP and optical network, with traditional approaches the planning and service provisioning would be complex and time consuming due of the manual synergy needed between the operator's IP team and optical team. With the introduction of SRv6 and the new SRv6 behaviors defined in this document, one simplified approach for IP and optical integration is to build a SRv6 SID list that integrates the path in both the IP layer and the optical layer.

As the optical layer is not packet based, source routing mechanism can not be directly used in the optical network. However, the abstracted optical paths (e.g., with ODUk or DWDM) could be exposed to the control system of the IP network using the SRv6 End.IL SIDs, and some of the attributes of the optical paths may also be provided. Based on this information, IP-optical inter-layer paths can be computed and programmed using SRv6 SID lists to meet some specific service requirements, such as low latency.

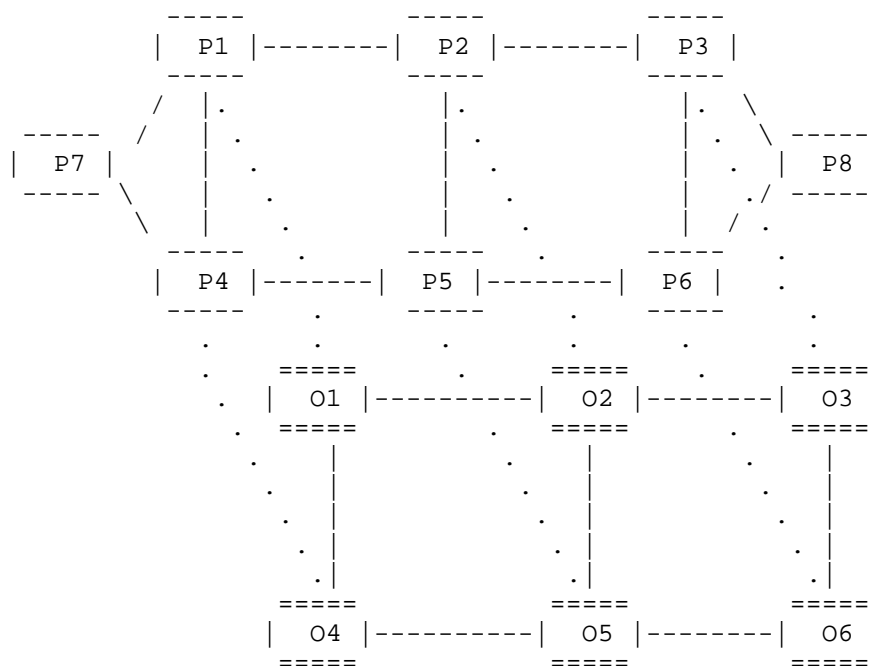


Figure 1. IP and Optical Layered Network Topology

In Figure 1, P1 to P8 are IP nodes, and O1 to O6 are optical nodes. Assume the operator needs to deploy a low latency path between P7 and P8. With normal segment routing, an IP layer path with the segment list {P7, P1, P2, P3, P8} can be used. If an optical path from O1 to O3 exists, the End.IL SID as defined in this document can be used to announce this optical path as an underlay interface or connection with specific attributes into the IP network. The headend node or the controller in IP layer can program an inter-layer TE path along {P7, P1, (O1, O2, O3), P3, P8} which could provide lower latency.

The underlay optical path between O1 and O3 may be created in advance or as a result of the request from the IP layer. The creation should be done by the optical network controller (not shown in the figure). The details of the process are out of scope of this document, and may refer to [I-D.ietf-teas-actn-poi-applicability].

There is also another case of IP and Optical integration. Assume there are two optical paths between P1 and P2. One is {P1, O1, O2, P2}, and the other is {P1, O1, O4, O5, O2, P2}. With SRv6 inter-layer programming, two separate SRv6 inter-layer SIDs can be allocated for these two underlay connections respectively. One is P1::C2 for the underlay path {P1, O1, O2, P2}, and the other is

P1::C45 for the path {P1, O1, O4, O5, O2, P2}. The headend P7 or the IP network controller will be informed about these two SRv6 SIDs and the associated path attributes, so that the headend or the controller can program different end-to-end inter-layer paths using SRv6 SID lists with different SRv6 inter-layer SIDs for services with different SLA requirements.

4.2. IP and MTN Integration

Assuming that an operator owns both an MTN network domain and an IP network domain. In the MTN network, each MTN node has both the layer-3 functionality and the MTN Path layer functionality. In layer-3, all the MTN nodes are in a layer-3 network topology, which connects to the IP network domain. In the MTN Path Layer, a set of MTN paths are provisioned between the selected pairs of MTN nodes for traffic engineering. In the MTN network, different types of services may be carried using either a layer-3 path, an end-to-end MTN path, or an inter-layer path comprising of both the layer-3 links and the MTN paths as segments. In addition, For some type of services, end-to-end paths across the IP domain and the MTN domain are needed, which is comprised of both the layer-3 paths and the MTN path as different segments.

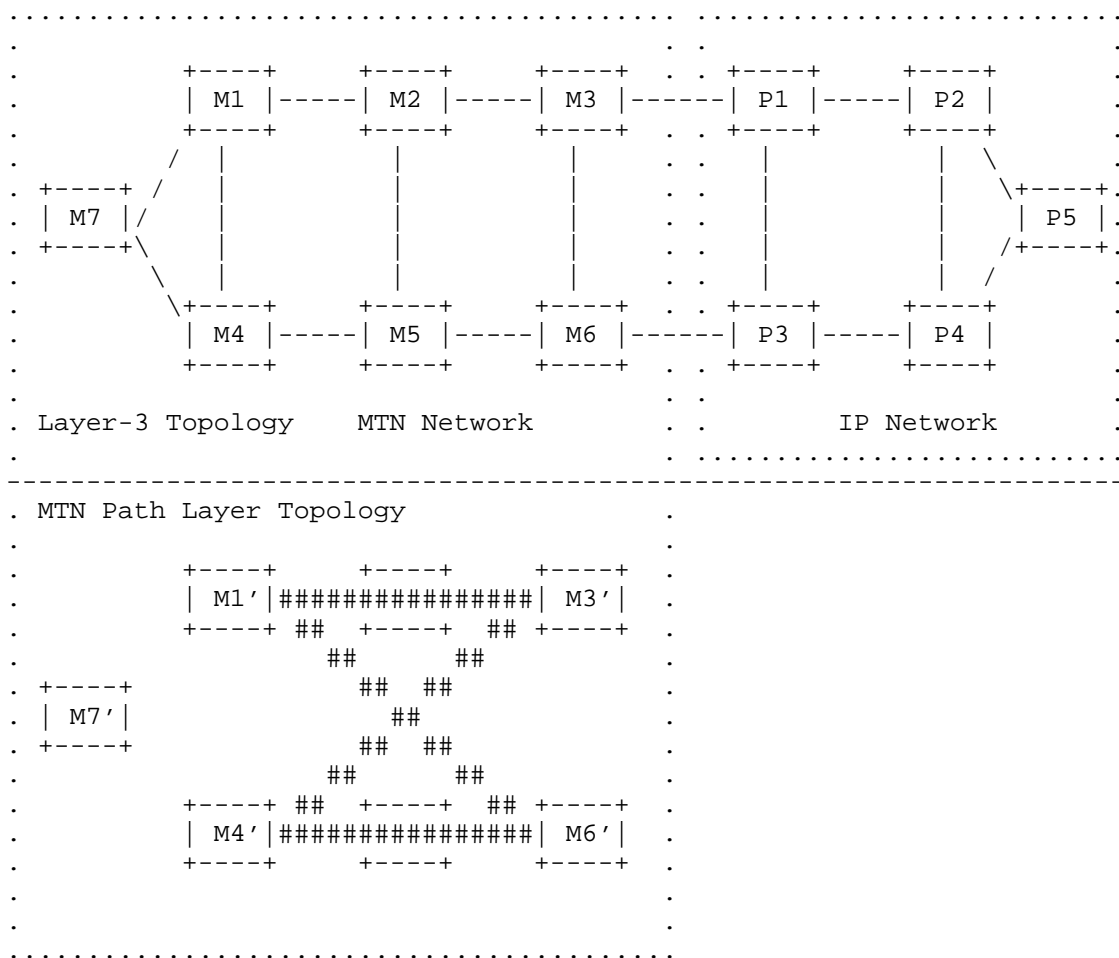


Figure 2. A network with MTN Domain and IP Domain

Figure 2 gives an example of a network with a MTN domain and an IP domain. M1 to M7 are MTN nodes, and P1 to P4 are IP nodes. The same set of MTN nodes builds two separate network layers. The topology in the IP layer shows the layer-3 connectivity between the MTN nodes and the connectivity with the IP network domain, while the topology in the MTN Path layer shows the MTN paths between the selected pair of MTN nodes. An end-to-end path from M7 to P5 can be established in layer-3 using an SRv6 SID list representing the layer-3 path {M7, M1, M2, M3, P1, P2, P5}. While for services which require low latency, with SRv6 inter-layer programming, an end-to-end path consisting of both the layer-3 segments and MTN paths could be established using an SRv6 SID list representing the inter-layer path {M7, M1::C3, P1, P2, P5}, where the SRv6 SID M1::C3 represents the underlay MTN path M1'-M3'.

This shows that it is convenient to use integrated SRv6 SID lists to program inter-layer TE paths both within the MTN domain, and across the IP and MTN domain using the combination of SRv6 L3 SIDs and the SRv6 inter-layer SID.

5. Security Considerations

The security considerations of SRv6 in [RFC8754] [RFC8986] apply to this document. As the information of underlay network connection is considered sensitive, operators need to make sure such information exposure only happen under the same administration, and proper policy is applied to ensure that only the exposure of necessary information is allowed.

6. IANA Considerations

This document defines a new SRv6 Endpoint behavior called END.IL.

IANA has allocated the following code points for different flavors of End.IL from the "SRv6 Endpoint Behaviors" sub-registry in the "Segment-routing with IPv6 data plane (SRv6) Parameters" registry:

Value	Hex	Endpoint Behavior	Reference
150	0x0096	End.IL	[This ID]
151	0x0097	End.IL with PSP	[This ID]
152	0x0098	End.IL with USP	[This ID]
153	0x0099	End.IL with USD	[This ID]
154	0x009A	End.IL with PSP, USP & USD	[This ID]
155	0x009B	End.IL with REPLACE-CSID	[This ID]
156	0x009C	End.IL with REPLACE-CSID & PSP	[This ID]
157	0x009D	End.IL with REPLACE-CSID, PSP, USP & USD	[This ID]

7. Acknowledgements

The authors would like to thank Xiaodong Chang, Yongjian Hu, Alexander Vainshtein, Ketan Talaulikar and Zhibo Hu for their review and comments.

8. References

8.1. 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>>.
- [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>>.

8.2. Informative References

- [I-D.ietf-teas-actn-poi-applicability] Peruzzini, F., Bouquier, J., Busi, I., King, D., and D. Ceccarelli, "Applicability of Abstraction and Control of

Traffic Engineered Networks (ACTN) to Packet Optical Integration (POI)", Work in Progress, Internet-Draft, draft-ietf-teas-actn-poi-applicability-14, 25 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-actn-poi-applicability-14>>.

[ITU-T_G.8310]

ITU-T, "ITU-T G.8310: Architecture of the metro transport network", <https://www.itu.int/rec/T-REC-G.8310-202012-I/en>, December 2020.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.

[RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <<https://www.rfc-editor.org/info/rfc6241>>.

[RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

Authors' Addresses

Liuyan Han
China Mobile
No.32 XuanWuMen West Street
Beijing, 100053
China
Email: hanliuyan@chinamobile.com

Jie Dong
Huawei Technologies
Huawei Campus, No.156 Beiqing Road
Beijing, 100095
China
Email: jie.dong@huawei.com

Minxue Wang
China Mobile
No.32 XuanWuMen West Street
Beijing, 100053
China
Email: wangminxue@chinamobile.com

Ran Chen
ZTE Corporation
Nanjing
China
Email: chen.ran@zte.com.cn

Zongpeng Du
China Mobile
No.32 XuanWuMen West Street
Beijing, 100053
China
Email: duzongpeng@foxmail.com