

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 6 May 2026

G. Dong
C. Xie
China Telecom
X. Li
CERNET Center/Tsinghua University
S. Peng
Huawei Technologies
E. Metz
KPN
2 November 2025

4map6 Segments for IPv4 Service delivery over IPv6-only underlay
networks
draft-dong-spring-sr-4map6-segments-05

Abstract

This document defines a new type of segment for Segment Routing, 4map6 segment, along with its corresponding End.4map6 endpoint behavior. This behavior performs stateless IPv4/IPv6 translation or encapsulation at PE nodes based on stateless mapping rules, thereby enabling the delivery of IPv4 services over an IPv6-only network. At the same time, the BGP Prefix-SID attribute SRv6 Service TLVs is extended to transmit IPv4/IPv6 address mapping rules in IPv6-only domain and cross-domain.

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 6 May 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. Overview of the new SID	3
3. SRv6 Service TLV Extension	4
4. Behavior	5
5. IANA Considerations	6
6. Security Considerations	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Authors' Addresses	8

1. Introduction

The document [I-D./draft-ietf-v6ops-framework-md-ipv6only-underlay] proposes a framework for deploying IPv6-only as underlay in multi-domain networks. In this framework, each PE will be identified by at least one IPv6 mapping prefix assigned by the operator, and it will be routable. It will also have one or more associated IPv4 address blocks extracted from the local IPv4 routing table or address pool. In addition, a specific data structure is defined as address mapping rules to express the mapping relationship between IPv4 address blocks and the IPv6 mapping prefixes of the remote PE. With this design, if the ingress PE obtains the mapping rule of a remote PE, the rule will provide forwarding guidance for IPv4 packets in the IPv6-only network when the destination address of the IPv4 packet matches its IPv4 address block. The ingress PE will use the information in the mapping rule to generate the corresponding IPv6 source and destination addresses from its IPv4 source and destination addresses and then generate a new IPv6 packet.

This mapping-based conversion can also work in SRv6 [RFC8986] networks. SRv6 defines packet processing in IPv6 networks as a list of instructions, which are represented as 128-bit segments, often referred to as Segment ID (SID). In this document, a new type of segments, 4map6 segments, is defined for Segment Routing. They run in PE nodes and provide support for implementing IPv4/IPv6 conversion

function based on mapping rules. In multi-domain IPv6-only networks, ingress nodes can convert IPv4 packets into IPv6 packets by stateless encapsulation or translation using the information in 4map6 segments, and egress nodes can restore IPv4 packets after transmission in IPv6-only networks. This draft illustrates IPv4/IPv6 address mapping and packet transformation from the perspective of SRv6. The new SID mainly deals with how to generate new IPv6 headers based on the information of IPv4 packets. It does not involve the SRH, so it can be considered as an approach of SRv6 BE.

The design of this draft only supports the scenarios in section 3 of the document [I-D./draft-ietf-v6ops-framework-md-ipv6only-underlay].

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. Overview of the new SID

The SRv6 SID has the same format as a 128-bit IPv6 address. According to Section 3 of [RFC8986], the SID consists of LOC:FUNCT:ARG, where the address (LOC) is encoded in the L most significant bits of the SID, followed by F bits of the function (FUNCT), and A bits of the argument (ARG).




Figure 1: SID architecture

The LOC identifies the node where the SID is instantiated and directs the packet to it. The FUNCT is an opaque identifier for a behavior bound to SID. The ARG field provides additional information for its processing. As a new type of SID, the 4map6 segment follows the format of the general SID. In addition, some information items specific to stateless address mapping and packet translation are carried in the relevant fields of the 4map6 SID as follows:

■ The LOC field has the positioning function, which is unique in the SR domain. It is a prefix assigned by the operator, and its value reflects the operator's IPv6 address planning. The operator needs to plan this prefix according to its own networking and business classification, and it is used to identify the node that instantiates the 4map6 SID.

「 The FUNCT field identifies the behavior associated with the 4map6 SID and instructs the node that instantiates the 4map6 SID node to perform the corresponding functional operation.

「 The ARG field contains the behavioral identity about whether stateless encapsulation or translation is performed at that point, as well as the IPv4 address associated with the PE node. The value of L+F should be less than or equal to 96 since 32 bits are required for IPv4 address.


For a PE node instantiating a 4map6 SID, its IPv6 mapping prefix IPv6 Prefix for IPv4 transmission corresponds to the LOC field. For the encapsulation or translation processing of IPv4 packets E/T is distinguished by the identity of ARG field. Therefore, the 4map6 SID has the structure: IPv6 Prefix + E/T + IPv4 address. In general, the number of SIDs that can be instantiated on a PE node is equal to the number of associated IPv4 address blocks.

3. SRv6 Service TLV Extension

In SRv6 network environment, 4map6 SID needs to be published to implement IPv4/IPv6 address mapping rule advertisement. This document specifies that an egress PE can implement the BGP protocol and extend in its BGP Prefix-SID property [RFC8669] to support publishing the service SID contained in its TLV, that is, the 4map6 SID, as an overlay service prefix in the network. The standard BGP update propagation scheme [RFC4271] can make use of route reflectors [RFC4456] to propagate these prefixes. This document defines a new BGP prefixed SID attribute extension TLV in the SRv6 Service TLVs to implement SID signaling for the 4map6 service of the SRv6 service:

SRv6 4map6 Service TLV: Used to carry SRv6 SID information of 4map6 service in multi-domain IPv6-only network, extended to support carrying End.4map6 SID.

The format of the extended SRv6 Services TLV is depicted below:



TLV Type (8 bits):

Used to identify different TLVS, SRv6 4map6 Service TLV is 8.

TLV Length (16 bits):
Total length of the TLV.

```
RESERVED (8 bits):
    Reserved fields.
```

SRv6 Service Sub-TLVs (variable):

This field contains more specific SRv6 service information and is encoded as an unordered list of Sub-TLVs whose format is described below.

The format of SRv6 Service Sub-TLVs is depicted below:

```

0 7 15 23 31
+-----+-----+-----+-----+ | SRv6
Service | SRv6 Service | SRv6 Service | | Sub-TLV Type | Sub-TLV
Length | Sub-TLV Value|
+-----+-----+-----+-----+
SRv6 Service Sub-TLV

```

Figure 3:

SRv6 Service Sub-TLV Type (8 bits):
This field identifies the type of SRv6 service information, which takes the value 1, indicates the SRv6 SID Information Sub-TLV.

SRv6 Service Sub-TLV Length (16 bits):
The sub-TLV length.

SRv6 Service Sub-TLV Value (variable):
This field contains data specific to the Sub-TLV Type. The SRv6 SID information of 4map6 Service is filled in the SRv6 Service Sub-TLV Value field.

4. Behavior

In general, 4map6 SID nodes operate in pairs. For a particular data stream, one node is the ingress PE, denoted by PE1, and the other is the egress PE, denoted by PE2. Each PE maintains a Mapping rule Database (MD) as shown in Figure. The table entries in the MD database consist of IPv4 address prefixes, IPv6 mapping prefixes, and the encapsulation or translation processing E/T way of IPv4 packets. It should be noted that the database here is only an example, and developers can design the structure of the database according to the actual situation.

IPv4 Address Prefix	IPv6 Mapping Prefix	Encapsulation or Translation E/T

Figure 4: The Structure of Map Rule Database

Before transmitting an IPv4 packet from PE1 to PE2, the address mapping rule corresponding to its IPv4 destination address needs to be transferred from PE2 to PE1. Therefore, PE2, which supports 4map6 service based on SRv6, publishes the 4map6 SID corresponding to the IPv4 destination address in the BGP Prefix-SID attribute, so as to realize the PE2 node located at the edge of the network to other nodes and synchronize the address mapping rule database on each PE device. PE2 announces its capabilities to other nodes in the format of the 4map6 SID of the SRv6 domain in the control plane. When PE1 receives the 4map6 SID announced by PE2, it extracts the relevant information and stores it in the local mapping rule database. The LOC field in the 4map6 SID corresponds to the database IPv6 Mapping Prefix, the Encapsulation/Translation behavior identifier bit in the ARG corresponds to the database encapsulation or translation E/T, and the address field in the ARG corresponds to the database IPv4 address block.

When PE1 receives an IPv4 packet, it first uses the destination IPv4 address block to find the corresponding IPv4 address block entry in the local mapping rule database. If there is a matching IPv4 address block entry, the corresponding IPv6 mapping prefix will concatenate the IPv4 destination address to generate the 4map6 SID, and the 32-bit IPv4 destination address in the packet is placed in the ARG field. According to the general SRv6 procedure, the SID is programmed as an IPv6 destination address, and the newly generated packet is sent to the pure IPv6 network for further delivery.

When a new IPv6 packet arrives at PE2, PE2 resolves the IPv6 destination address of the packet. If it matches an IPv6 mapping prefix in its instantiated SID, it will process the packet according to the 4map6 instruction in FUNCT, remove the IPv6 mapping prefix and forward it to the next hop according to the encapsulation/translation behavior in ARG field and the destination IPv4 address carried.

5. IANA Considerations

This document requests IANA to allocate the following codepoints in "SRv6 Endpoint Behaviors" sub-registry of "Segment Routing Parameters" top-level registry.

Value	Hex	Endpoint behavior	Reference
TBD		End.4map6	[This.ID]

Figure 5: End.4map6 Endpoint Behavior

6. Security Considerations

There is no additional security risk introduced by this design.

7. References

7.1. 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>>.
- [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>>.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<https://www.rfc-editor.org/info/rfc4456>>.
- [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] Filts, 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>>.
- [RFC8669] Previdi, S., Filts, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <<https://www.rfc-editor.org/info/rfc8669>>.
- [RFC8754] Filts, 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] Filts, 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>>.

[RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <<https://www.rfc-editor.org/info/rfc9252>>.

7.2. Informative References

[I-D.ietf-v6ops-framework-md-ipv6only-underlay] Xie, C., Ma, C., Li, X., Mishra, G. S., and T. Graf, "Framework of Multi-domain IPv6-only Underlay Network and IPv4-as-a-Service", Work in Progress, Internet-Draft, draft-ietf-v6ops-framework-md-ipv6only-underlay-14, 8 October 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-framework-md-ipv6only-underlay-14>>.

Authors' Addresses

Guozhen Dong
China Telecom
Beiqijia Town, Changping District
Beijing
102209
China
Email: donggz@chinatelecom.cn

Chongfeng Xie
China Telecom
Beiqijia Town, Changping District
Beijing
102209
China
Email: xiechf@chinatelecom.cn

Xing Li
CERNET Center/Tsinghua University
Shuangqing Road No.30, Haidian District
Beijing
100084
China
Email: xing@cernet.edu.cn

Shuping Peng
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.

Beijing
100095
China
Email: pengshuping@huawei.com

Eduard Metz
KPN
Netherlands
Email: etmetz@gmail.com