Network C. Weiqiang
Internet-Draft China Mobile
Intended status: Standards Track G. Mirsky
Expires: May 6, 2020 ZTE Corp.
P. Shaofu
L. Aihua
ZTE Corporation
W. Xiaolan
New H3C Technologies Co. Ltd
C. Wei
Centec
S. Zadok
Broadcom
November 3, 2019

Unified Identifier in IPv6 Segment Routing Networks
draft-mirsky-6man-unified-id-sr-04

Abstract

Segment Routing architecture leverages the paradigm of source routing. It can be realized in a network data plane by prepending the packet with a list of instructions, a.k.a. segments. A segment can be encoded as a Multi-Protocol Label Switching (MPLS) label, IPv4 address, or IPv6 address. Segment Routing can be applied in MPLS data plane by encoding segments in MPLS label stack. It also can be applied to IPv6 data plane by encoding a list of segment identifiers in IPv6 Segment Routing Extension Header (SRH). This document extends the use of the SRH to unified identifiers encoded as MPLS label or IPv4 address, to compress the SRH, and support support more detailed network programming and interworking between SR-MPLS and SRv6 domains.

This Internet-Draft will expire on May 6, 2020.

Copyright Notice

Table of Contents

1. Introduction

Segment Routing architecture [RFC8402] leverages the paradigm of source routing. It can be realized in a network data plane by prepending the packet with a list of instructions, a.k.a. segment identifiers (SIDs). A segment can be encoded as a Multi-Protocol Label Switching (MPLS) label, IPv4 address, or IPv6 address. Segment Routing can be applied in MPLS data plane by encoding 20-bits SIDs in MPLS label stack [I-D.ietf-spring-segment-routing-mpls]. It also can be applied to IPv6 data plane by encoding a list of 128-bits SIDs in IPv6 Segment Routing Extension Header (SRH)

[I-D.ietf-6man-segment-routing-header]. Applicability of 32-bits SID that may represent an IPv4 address has not been defined.

SR extensions to Interior Gateway Protocols (IGP), IS-IS [I-D.ietf-isis-segment-routing-extensions], OSPF [I-D.ietf-ospf-segment-routing-extensions], and OSPFv3 [I-D.ietf-ospf-ospfv3-segment-routing-extensions], defined how 20-bits and 32-bits SIDs advertised and bound to SR objects and/or instructions. Extensions to BGP link-state address family [I-D.ietf-idr-bgp-ls-segment-routing-ext] enabled propagation of segment information of variable length via BGP.

This document extends the use of the SRH [I-D.ietf-6man-segment-routing-header] to unified identifiers encoded as MPLS label or IPv4 address to support more detailed network programming and interworking between SR-MPLS and SRv6 domains.

1.1. Conventions used in this document

1.1.1. Terminology

SR: Segment Routing

SRH: Segment Routing Extension Header

MPLS: Multiprotocol Label Switching

SR-MPLS: Segment Routing using MPLS data plane

SID: Segment Identifier

IGP: Interior Gateway Protocol

DA: Destination Address

ILM: Incoming Label Map

FEC: Forwarding Equivalence Class

FTN: FEC-to-NHLFE map

OAM: Operation, Administration and Maintenance

TE: Traffic Engineering

SRv6: Segment Routing in IPv6

U-SID: Unified Segment Identifier

PSP: Penultimate Segment Popping

FIB: Forwarding Information Base

1.1.2. 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. Segment Routing Extension Header: Benefits and Challenges

Many functions related to Operation, Administration and Maintenance (OAM) require identification of the SR tunnel ingress and the path, constructed by segments, between the ingress and the egress SR nodes. Combination of IPv6 encapsulation [RFC8200] and SRH [I-D.ietf-6man-segment-routing-header], referred to as SRv6, comply with these requirements while it is challenging when applying SR in MPLS networks, also referred to as SR-MPLS.

On the other hand, the size of IPv6 SID presents a scaling challenge to use topological instructions that define strict explicit traffic engineered (TE) path or support network programming in combination with service-based instructions. At the same time, that is where SR-MPLS approach provides better results due to smaller SID length. It can be used to compress the SRv6 header size when a smaller namespace of available SIDs is sufficient for addressing the particular network.

SR-MPLS is broadly used in metro networks. With the gradual deployment of SRv6 in the core networks, supporting interworking between SR-MPLS and SRv6 becomes the necessity for operators. It is operationally more efficient and straightforward if SRv6 can use the same size SIDs as in SR-MPLS. The SRH can be extended to define the same as in SR-MPLS SID length to support the unified segment identifier (U-SID). As a result, end-to-end SR tunnel may use U-SIDs across SR-MPLS and SRv6 domains.

3. Unified SIDs in IPv6 Segment Routing Extension Header

SRH format has been defined in Section 3 of [I-D.ietf-6man-segment-routing-header] as presented in Figure 1

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    | Next Header   |  Hdr Ext Len  | Routing Type  | Segments Left |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |  Last Entry   |     Flags     |              Tag              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    |            Segment List[0] (128 bits IPv6 address)            |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    |                                                               |
                                  ...
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    |            Segment List[n] (128 bits IPv6 address)            |
    |                                                               |
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    //                                                             //
    //         Optional Type Length Value objects (variable)       //
    //                                                             //
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1: SRH format

This document defines a new field Size in the SRH Flags field as a two-bits field with the following values:

0b00 - 128-bits SID, an IPv6 address;

0b01 - 32-bits SID, an IPv4 address;

0b10 - 32-bits SID, an MPLS label in leftmost 20-bits, rightmost 12-bits for context information used by the label forwarding entry. The context information could be U-SID function code.

0b11 - reserved for future use.

Entries of the segment list in the SRH MUST be of the same length.

4. The Use Case of Unified Segment Identifier

U-SID can be used for interworking between SR-MPLS and SRv6 domains. SR-MPLS is often used in a metro network, for example, in the backhaul metro network of CMCC. If the core network uses SRv6, for example, the core network of the same operator, U-SID can be used in the SRv6 domain to interwork with SR-MPLS in the metro network to form an end-to-end tunnel.

4.1. Interworking Between SR-MPLS and SRv6 Using U-SID

SR-MPLS uses SR SIDs as MPLS label in MPLS stack, and the SIDs are 32-bits long. SRv6 uses SR SIDs as IPv6 extension header in SRH, and the SIDs are 128-bits long.

The U-SID uses the same 32-bits long SIDs in MPLS stack and SRH. Thus, four 32-bits long U-SIDs can be placed in the space of a single 128-bits long header. The encapsulation is illustrated in Figure 2.

```
        +---------+          +-----------------------------------+
        |         |          |           IPv6 header             |
        | Ethernet|          +-----------------------------------+
        |         |          |           SRH                     |
        +---------+          +-----------------------------------+
        |  USID1  |          | USID1  | USID2  | ...   | USID4   |
        +---------+          +-----------------------------------+
        |  USID2  |          | USID5  |...     | USIDn | Null    |
        +---------+          +-----------------------------------+
        | ...     |          +           Payload                 |
        +---------+          +-----------------------------------+
        |  USIDn  |
        +---------+
        | Payload |
        +---------+
```

Figure 2: 32-bits long U-SIDs Encapsulation

The SR-MPLS and SRv6 interworking is illustrated in Figure 3. An end-to-end SR tunnel from A to F crosses the SR-MPLS and SRv6 domains. The SR-MPLS domain could be using IPv4 or IPv6 address family. The SRv6 border nodes (E/G) receive SR-MPLS packets and forward them into the SRv6 domain using an SR-MPLS Binding SID [I-D.ietf-spring-segment-routing-mpls].

```
     +-----+           +-----+           +-----+           +-----+
     |  A  +-----------+  B  +-----------+  E  +-----------+  F  |
     +-----+           +--+--+           +--+--+           +--+--+
        |    SR-MPLS      |                 |     SRv6        |
        |                 |                 |                 |
     +-----+           +--+--+           +--+--+           +--+--+
     |  C  |-----------|  D  +-----------+  G  +-----------+  H  |
     +-----+           +-----+           +-----+           +-----+

                                                +--------------+
                                                |   Eth(E->G)  |
        +--------------+                        +--------------+
        |   Eth(A->B)  |                        |IPv6 DA:G.intf|
        +--------------+    +--------------+    +--------------+
        |   USID(B)    |    |   Eth(B->E)  |    |SRH           |
        +--------------+    +--------------+    |NH:MPLS   SL:2|
        |   USID(E1)   |    |   USID(E1)   |    |USID(ADJ E->G)|
        +--------------+    +--------------+    |USID(ADJ G->H)|
        |   USID(E2)   |    |   USID(E2)   |    |USID(ADJ H->F)|
        +--------------+    +--------------+    +--------------+
        |   USID(F)    |    |   USID(F)    |    |   USID(F)    |
        +--------------+    +--------------+    +--------------+
        |Label(service)|    |Label(service)|    |Label(service)|
        +--------------+    +--------------+    +--------------+
        |    Payload   | -> |    Payload   | -> |    Payload   |
        +--------------+    +--------------+    +--------------+
```

Figure 3: SR-MPLS and SRv6 interworking

The SRv6 edge node E assigns two SIDs, e.g., E1 and E2, E1 is an SR-MPLS Node-SID, E2 is an SR-MPLS Binding-SID, which represents an SRv6 policy (from E to F, via segment list E-G-H-F) with U-SID encapsulation. Figure 4 demonstrates an example of the packet forwarding, where U-SID is an MPLS label.

The controller may assign the end-to-end SR tunnel U-SIDs (from A to F), and another method is outside the scope of this document.

5. Operations with Unified Segment Identifier

When SRH is used to include 32-bits long U-SIDs, the ingress and transit nodes of an SR tunnel act as described in Section 5.1 and Section 5.2 of [I-D.ietf-6man-segment-routing-header] respectively.

If U-SID is used to support interworking between SR-MPLS and SRv6 domains, it is beneficila that U-SID type matches to an MPLS label. In that case, an ILM (Incoming Label Map) entry can be used to map a

U-SID to an IPv6 address. As result, it is not necessary to introduce a new type of index-based mapping table. For ILM entry of Adjacency-SID, the mapping result copied to DA (Destination Address) is the remote interface IPv6 address, for ILM entry of Node-SID, the mapping result copied to DA is remote node loopback IPv6 address.

Operations oon an MPLS label of U-SID type are the same as those defined in [I-D.ietf-mpls-sr-over-ip]. However, SR-MPLS over SRH has the following advantages compared with SR-MPLS over UDP:

- o SRH is flexible to extend flags or sub-TLVs for service requirements, but UDP not.
- o Labels in SRH can meet 8 bytes alignment requirements as per [RFC8200], but UDP not.
- o The source address of the SR policy is not discarded, but UDP not.

5.1. Procedures of SR-MPLS over IP

Procedures of SR-MPLS over IP of [I-D.ietf-mpls-sr-over-ip] described how to construct an adjusted SR-MPLS FTN (FEC-to-NHLFE map) and ILM entry towards a prefix-SID when next-hops are IP-only routers, the action of FTN and ILM entry will steer the packet along an outer tunnel to the target node that originated the FEC (Forwarding Equivalence Class), and on each airway node along the segment list, UDP header is frequently removed and put again. However, for SR-MPLS over SRH in this document we don't try to depend on that adjusted FIB (Forwarding Information Base) entry, because there are not any actions needed to get from the FIB entry, a traditional ILM entry (maybe without out-label because of IP-only next-hop) is enough to get the FEC information, i.e., to map a U-SID to an IPv6 address and copy to DA. An SRv6 policy chosen to encapsulate U-SID list within SRH is determined at the ingress node of this SRv6 policy, SRH is preserved along the SR to egress, though PSP (Penutimate Segment Popping) may be used, that is different from SR-MPLS over IP/UDP method [I-D.ietf-mpls-sr-over-ip], so the source address (i.e., the ingress of the SRv6 policy) is not discarded.

5.2. Packet Forwarding

U-SID based packet forwarding is similar to the processing described in [I-D.ietf-mpls-sr-over-ip]. But it differs from that in FIB action and segment list processing. For completeness, we repeat the description of [I-D.ietf-mpls-sr-over-ip] with modification as follows.

```
      +-----+       +-----+       +-----+        +-----+        +-----+
      |  A  +-------+  B  +-------+  C  +--------+  D  +--------+  H  |
      +-----+       +--+--+       +--+--+        +--+--+        +-----+
                       |             |              |
                       |             |              |
                    +--+--+       +--+--+        +--+--+
                    |  E  +-------+  F  +--------+  G  |
                    +-----+       +-----+        +-----+

           +--------+          +--------+          +--------+
           |IP(A->E)|          |IP(A->G)|          |IP(A->G)|
           +--------+          +--------+          +--------+
           |SRH     |          |SRH     |          |SRH     |(or PSP)
           |  SL:2  |          |  SL:1  |          |  SL:0  |
           |  L(E)  |          |  L(E)  |          |  L(E)  |
           |  L(G)  |          |  L(G)  |          |  L(G)  |
           |  L(H)  |          |  L(H)  |          |  L(H)  |
           +--------+          +--------+          +--------+
           | Packet |   --->   | Packet |     ---> | Packet |
           +--------+          +--------+          +--------+
```

Figure 4: Packet Forwarding Example

In the example shown in Figure 4, assume that routers A, E, G, and H are U-SID capable (i.e, both SR-MPLS and SRv6 capable) while the remaining routers (B, C, D, and F) are only capable of forwarding IP packets. Routers A, E, G, and H advertise their Segment Routing related information via IS-IS or OSPF.

Now assume that router A (the Domain ingress) wants to send a packet to router H (the Domain egress) via an SRv6 policy with the explicit path {E->G->H}. Router A will impose an MPLS label stack within SRH on the packet that corresponds to that explicit path. Router A searches ILM entry by the top label (that indicated router E), get the FEC information, a loopback IPv6 address of E, and then copy to DA and sends the packet. The value of SRH.SL is 2.

When the IPv6 packet arrives at router E, router E get the next segment (label) within SRH according to SL 2, searches ILM entry by the next label, get the FEC information, a loopback IPv6 address of G, and then copy to DA and sends the packet. The value of SRH.SL is 1.

When the IPv6 packet arrives at router G, router G gets the next segment (label) within SRH according to SRH.SL 1, looks up ILM entry by the next label, gets the FEC information, a loopback IPv6 address of H, and then copies it to IP DA and transmits the packet. Because

the value of SRH.SL is 0, the SRH can be removed if the Prefix-SID of H is set to PSP.

6. IANA Considerations

IANA is requested to allocate from the Segment Routing Header Flags registry the two-bits long field referred to as Size.

7. Security Considerations

This specification inherits all security considerations of [RFC8402] and [I-D.ietf-6man-segment-routing-header].

8. Acknowledgements

TBD

9. Normative References

[I-D.ietf-6man-segment-routing-header]
Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header (SRH)", draft-ietf-6man-segment-routing-header-26 (work in progress), October 2019.

[I-D.ietf-idr-bgp-ls-segment-routing-ext]
Previdi, S., Talaulikar, K., Filsfils, C., Gredler, H., and M. Chen, "BGP Link-State extensions for Segment Routing", draft-ietf-idr-bgp-ls-segment-routing-ext-16 (work in progress), June 2019.

[I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-25 (work in progress), May 2019.

[I-D.ietf-mpls-sr-over-ip]
Xu, X., Bryant, S., Farrel, A., Hassan, S., Henderickx, W., and Z. Li, "SR-MPLS over IP", draft-ietf-mpls-sr-over-ip-07 (work in progress), June 2019.

[I-D.ietf-ospf-ospfv3-segment-routing-extensions]
Psenak, P. and S. Previdi, "OSPFv3 Extensions for Segment Routing", draft-ietf-ospf-ospfv3-segment-routing-extensions-23 (work in progress), January 2019.

[I-D.ietf-ospf-segment-routing-extensions]
Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", draft-ietf-ospf-segment-routing-extensions-27 (work in progress), December 2018.

[I-D.ietf-spring-segment-routing-mpls]
Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with MPLS data plane", draft-ietf-spring-segment-routing-mpls-22 (work in progress), May 2019.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[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>.

Authors' Addresses

Cheng Weiqiang
China Mobile
Beijing
China

Email: chengweiqiang@chinamobile.com

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com

Peng Shaofu
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing
China

Email: peng.shaofu@zte.com.cn

Liu Aihua
ZTE Corporation
Zhongxing Industrial Park, Nanshan District
Shenzhen
China

Email: liu.aihua@zte.com.cn

Wan Xiaolan
New H3C Technologies Co. Ltd
No.8, Yongjia Road, Haidian District
Beijing
China

Email: wxlan@h3c.com

Cheng Wei
Centec
Building B, No.5 Xing Han Street, Suzhou Industrial Park
Suzhou
China

Email: Chengw@centecnetworks.com

Shay
Broadcom
Israel

Email: shay.zadok@broadcom.com