Authors: X. Geng   Z. Li
         Huawei Technologies   Huawei Technologies
         J. Xie
         Huawei Technologies

IPv6 Multicast Source Routing Traffic Engineering

Abstract

This document defines 2 new types of segment: End.RL and End.RL.X , and the corresponding packet processing procedures over the IPv6 data plane for the MSR6(Multicast Source Routing over IPv6) TE solutions.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].Ro

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Segment Routing ([RFC8402]) leverages the mechanism of source
routing. An ingress node steers a packet through an ordered list of
instructions, called "segments". Each one of these instructions
represents a function to be implemented at a specific location in
the network. A function is locally defined on the node where it is
executed. Network Programming combines Segment Routing functions to
achieve a networking objective that goes beyond mere packet routing.
[RFC8986] defines the SRv6 Network Programming concept and specifies
the main Segment Routing behaviors and network programming
functions.

Bit Index Explicit Replication (BIER) [RFC8279] is an architecture
that provides optimal multicast forwarding without requiring a
protocol for explicitly building multicast distribution trees or
per-flow state maintained by intermediate routers. When a multicast
data packet enters BIER forwarding domain, the ingress node
encapsulates the packet with a bitstring, each bitposition of which

presents the egress nodes. To forward the packet to a given set of egress nodes, the bits corresponding to those egress nodes are set in the bitstring. The intermediate nodes in the BIER domain replicate and forward the packet based on the bitstring.The mechanism of forwarding a packet based on bitstring of BIER are specified in [RFC8279].

An IPv6 based multicast source routing (MSR6) solution is defined in [I-D.cheng-spring-ipv6-msr-design-consideration], which leverages the benefits of source routing over IPv6 data plane to provide simplified multicast TE and BE service in an IPv6 network without unnecessary multicast tree status and complex control plane protocols. MSR6 needs to reuse the advantages of SRv6 and BIER to implement source routing.

MSR6 has two basic modes of forwarding: one is based on Shortest Path First(SPF), which is called MSR6 BE mode; the other is based on traffic engineered, which is called MSR6 TE mode. This document defines 2 new types of segment: End.RL and End.RL.X , and the corresponding packet processing procedures over the IPv6 data plane for the MSR6 TE solutions.

2. Terminologies

MSR6: Multicast Source Routing over IPv6, defined in [I-D.cheng-spring-ipv6-msr-design-consideration].

MRH: Multicast Routing Header, a new type of Routing Header which is used for MSR6.

Multicast domain: A set of network device which could provide P2MP multicast transport. In this document, the multicast domain is an MSR domain, where every nodes support the capability of MSR6.

Root Node: Root node is the beginning point of a multicast tree and also the initiation node of a multicast tunnel. It encapsulates the packet with a multicast header. The type of the encapsulation depends on the routing protocol used in the multicast domain. For MSR6 TE, the encapsulation is MSR6 TE header, which is an IPv6 header with MRH.

Replication Endpoint: the intermediate node of a multicast tree, which replicates packet and forwards the packet to the downstream nodes. For MSR6, the Replication Node is called Replication Endpoint which can be indicated by the MSR6 Segment and replicate packets according to the multicast source routing information encapsulation in the MSR6 header of the packet.

Leaf Node: Leaf node is the end point of a multicast tree and also the decapsulation node of a multicast tunnel. It decapsulates the

multicast header in the packet and forwards the packet based on overlay encapsulation.

Parent Node: The parent node is the node that does the packet replication, corresponding to the concept of a child node.

Child Node: The child node is the downstream node that will receive the packet which has been replicated ,corresponding to the concept of a parent node.

3. **MSR6 Explicit Routing**

In order to implement MSR6 TE, the nodes and links along the path must be explicitly specified in the packet.

In SRv6, the segment list is suitable for source route of unicast path, which is P2P and could be presented by the data structure of a chain with a pointer; In the Segment Routing Header(SRH), the segment list is the chain, and the segment left is the pointer. But for the multicast, the path of multicast service is a tree, and can not be presented by data structure in a one-dimensional array. So this document introduces new segment to present node/link in a multicast tree and introduce structure information in the segment to indicate the parent-child relationship during multicast replication. And all the nodes/links along the multicast tree can be encoded into the segment list.

In a multicast tree, when the packet is replicated from one node to multiple downstream nodes, parent-child relationship is built up between these nodes. MRH is supposed to encode the nodes along the P2MP path and also the parent-child relationship between them.

4. **MSR6 SID Definition**

As defined in [RFC8402], an Segment Identifier is an IPv6 address explicitly associated with the segment. When an SRv6 SID is in the Destination Address field of an IPv6 header of a packet, it is routed through transit nodes in an IPv6 network as an IPv6 address.

MSR6 Segment Identifier complies with the definition of the Segment Identifier in [RFC8402]. Following the specification in RFC8986, an MSR6 SID is consisting of LOC:FUNCT:ARG. Locator could be represented as B:N, B is the MSR6 segment block(IPv6 prefix allocated for MSR6 segment by the operator) and N is the identifier of the node instantiating the SID. The MSR6 locator is routable and leads to the node which instantiates the SID; Function is the identification of a local behavior bound to the MSR6 SID; Argument is the additional information requested by the function.

Typical MSR6 functions and arguments are defined in the following
sections.

5. MSR6 Endpoint Behaviors

5.1. End.RL: MSR6 Endpoint Replication List

End.RL SID is one of the basic MSR6 SID as an extension of End SID
defined in RFC8986.

The encoding of segment list which is composed of End.RL SIDs follow
the rules:

  *There MUST be an End.RL SID for each Replication Endpoint in the
   multicast tree

  *The End.RL SIDs, which represent child nodes that have the same
   parent node, MUST be arranged consecutively in the segment list

  *The End.RL SID for a parent node MUST indicate the position of
   the End.RL SIDs for the child nodes

There are 2 arguments for each End.RL SID:

  *"Replication number" is used for indicating how many times the
   parent node is supposed to execute replication

  *"Pointer" is used for indicating the position of End.RL SID for
   the first child node in the segment list;

When the packet is received by an Replication Endpoint and the DA of
this packet is a local SID with the function of End.RL, the packet
will be replicated based on the "replication number". The DA of the
1st replicated packet is replaced by the SL="pointer" and the value
of Segment Left is set to "pointer"; the DA of the 2nd replicated
packet is replaced by the SL="pointer +1" and the value of Segment
Left is set to "pointer+1". The packet is sent out based on the
updated DA. Repeat the operation untill all the replicated packets
are sent out.

The replication node does the following when the DA of the packet is
a local End.RL SID.

```
S01. When an MRH is processed {
S02.    If (Segments Left == 0 or Replication Number==Pointer==0) {
S03.        Stop processing the MRH, 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.    Replicate the packet based on the "replication number" in the
            argument of the SID
S14.    Set the Segment Left of the 1st replicated packet to "Pointer"
            the argument of the SID
S15.    Update IPv6 DA with Segment List[Pointer]
            respectively in each replicated packet
S16.    Submit the packet to the egress IPv6 FIB lookup for
            transmission to the new destination
S17.    Repeat S14-S17 with "Pointer+n"(n=the number of packets which
            been updated) untill all the packets are transmitted to base
            on the updated DA
R18.    }
```

5.2. End.RL.X: MSR6 L3 Cross-Connect

End.RL.X SID is one of the basic MSR6 SID as an extension of End.X
SID defined in RFC8986.

The encoding of segment list which is composed of End.RL.X SIDs
follow the rules:

  *There MUST be an End.RL.X SID for each downstream link connected
   to a Replication Endpoint in the multicast tree

  *The End.RL.X SIDs corresponding to downstream links of the same
   Replication Endpoint MUST be arranged consecutively in the
   segment list

  *The End.RL.X SID for a parent node MUST indicate the position of
   the End.RL.X SIDs for the child nodes

There are 2 arguments for each End.RL SID:

  *"Replication number" is used for indicating how many times the
   existing node is supposed to execute replication

  *"Pointer" is used for indicating the position of End.RL.X SID for
   the first downstream link of the child node in the segment list;

When the packet is received by a Replication Endpoint and the DA of
this packet is a local SID with the function of End.RL.X, the packet
will be replicated based on the "replication number". The DA of the
1st replicated packet is replaced by the Segment List[pointer] and
the value of Segment Left is set to "pointer"; The packet is sent
out through the link indicated by the SID. the DA of the 2nd
replicated packet is replaced by the Segment List [SL +1] and the
value of Segment Left is set to "SL+1"; The DA of the 2nd packet is
replaced by the Segment List [pointer2] (pointer2 is the argument
in the Segment List [SL+1]) and the value of Segment Left is set to
"pointer2"; The packet is sent out through the link indicated by the
SID; Repeat the operation untill all the replicated packets are sent
out through the indicated links.

```
S01. When an MRH is processed {
S02.    If (Segments Left == 0 or Replication Number==Pointer==0) {
S03.       Stop processing the MRH, 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.    Replicate the packet based on the "replication number" in the
            argument of the SID
S14.    Set the Segment Left of the 1st replicated packet to "Pointer"
            the argument of the SID
S15.    Update IPv6 DA with Segment List[Pointer]
S16.    Transmit the packet through the indicated link
S17.    Set the Segment Left of the 2nd replicated packet to "Segment
            the argument of the SID
S18.    Update IPv6 DA with Segment List[Segment Left+1]
S19.    Set the Segment Left of the 2nd replicated packet to "Pointer
            the argument of the updated SID
S20.    Update IPv6 DA with Segment List[Pointer 2]
S21.    Transmit the packet through the indicated link
S22.    Repeat S17-S21 with "Pointer+n"(n=the number of packets which
            been updated) untill all the packets are transmitted to base
            on the updated DA
R23.    }
```

6. Multicast Routing Header

A new type of Routing Header is defined as follows for MSR6 Traffic
Engineering called Multicast Routing Header(MRH).

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Next Header  |  Hdr Ext Len  |  Routing Type | Segments Left |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  MRH Type=1   |                   Reserved                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                           M-SID[0]                            |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                           M-SID[1]                            |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                             ...                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                           M-SID[n]                            |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Each M-SID in the MRH is shown in the following figure:

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                           Locator                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Function                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Replication number        |            Pointer            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
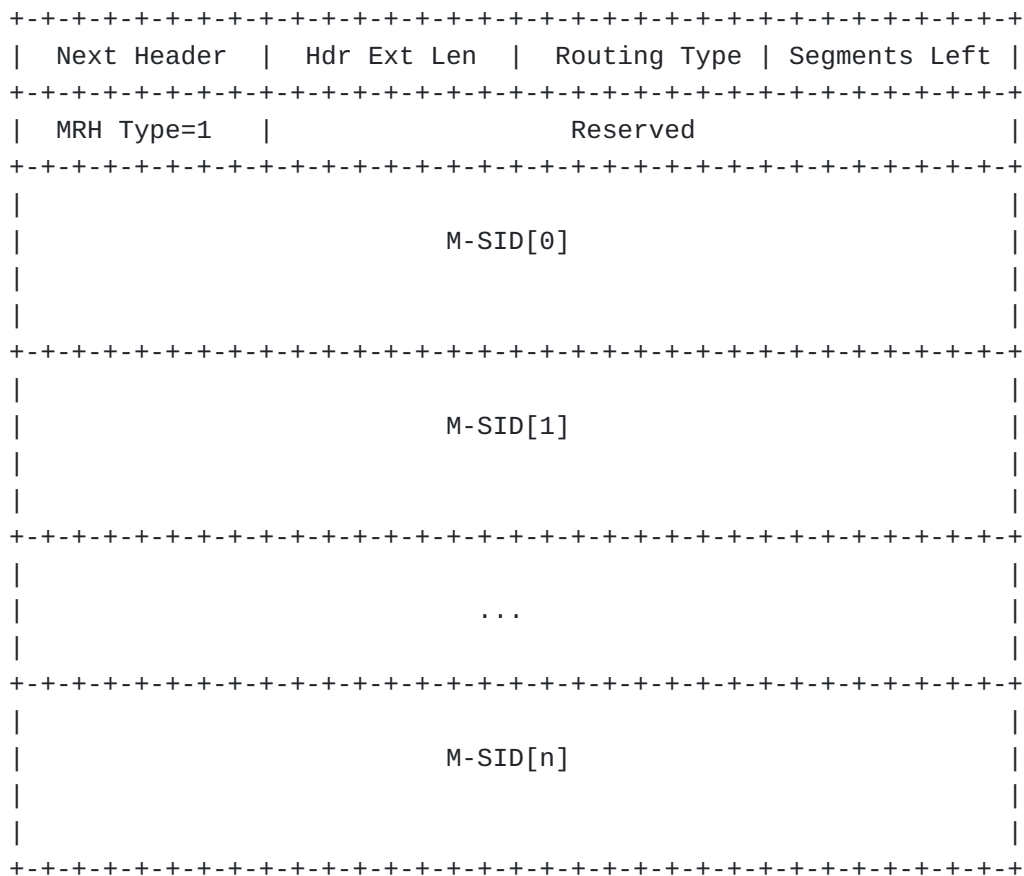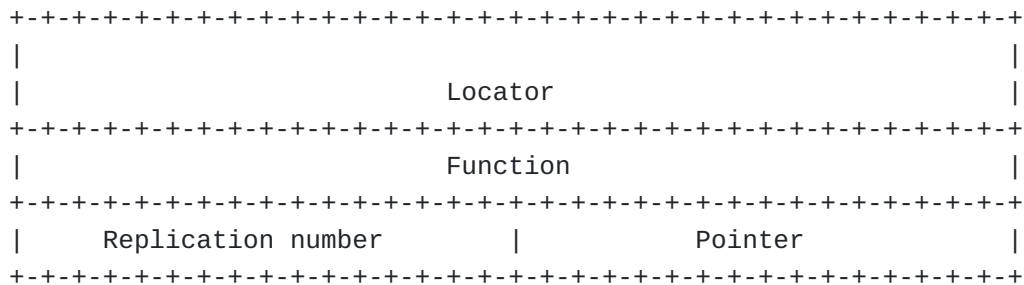
It includes:

Locator: could represented as B:N where B is the MRH segment block
(IPv6 prefix allocated for MRH segment by the operator) and N is the
identifier of the node instantiating the SID. When the locator part
of the MRH segment is routable, it leads to the node, which
instantiates the SID.

Function: an identification of a local behavior bound to the MRH.
Replication is default behavior for any MRH segment, which doesn't
need function indication.

Replication number is used for indicating how many times the existing node is supposed to excute replication

Pointer is used for indicating the position of MSR6 SID for the child node;

Segments for the group of child nodes which belong to the same parent node MUST be encoded together in the segment list. So Pointer and replication number could determine the upper bound and lower bound of the value range of segment left for the child nodes.

The destination address of the packet is the IPv6 address of the existing Replication Endpoint. The next header of the IPv6 header points to a Routing Header and the type of the routing header is MRH Type.

7. MSR6 Compression

7.1. MSR6 Requirement for Compression

Different from unicast, MSR is designed for an explicit multicast tree rather than a p2p path: the former contains more nodes than the latter generally. So the header overhead is one of the key challenge for MSR6, which limits the scalability and the potential deployment scenario.

There are some SRH compression solutions under discussion in IETF. These solutions could also be used for MSR6 with modification.

7.2. Compressed MRH6

The compressed MRH6 is designed as follows:

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Version  | Traffic Class |           Flow Label            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Payload Length       |   Next=43   |  Hop  Limit   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                     Source Address                            |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                     Common Prefix                             |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       M-SID [x]                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Next Header  |  Hdr Ext Len  |  Routing Type | Segments Left |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  MRH Type=2   |                  Reserved                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       M-SID[0]                                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       M-SID[1]                                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          ...                                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       M-SID[n]                                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Each M-SID in the MRH is shown in the following figure:

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Node-ID            | Function  | Rp Nm |Pointer|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

It includes:

Node ID: is the identifier of the node instantiating the SID. Common Profix and Node ID together form a routable IPv6 address

Function: an identification of a local behavior bound to the MRH. Replication is default behavior for any MRH segment, which doesn't need function indication.

Arguments: follow the definition in section 6, including replication number and pointer.

The processing process is:

The destination address of the packet is the IPv6 address of the
existing Replication Endpoint. The next header of the IPv6 header
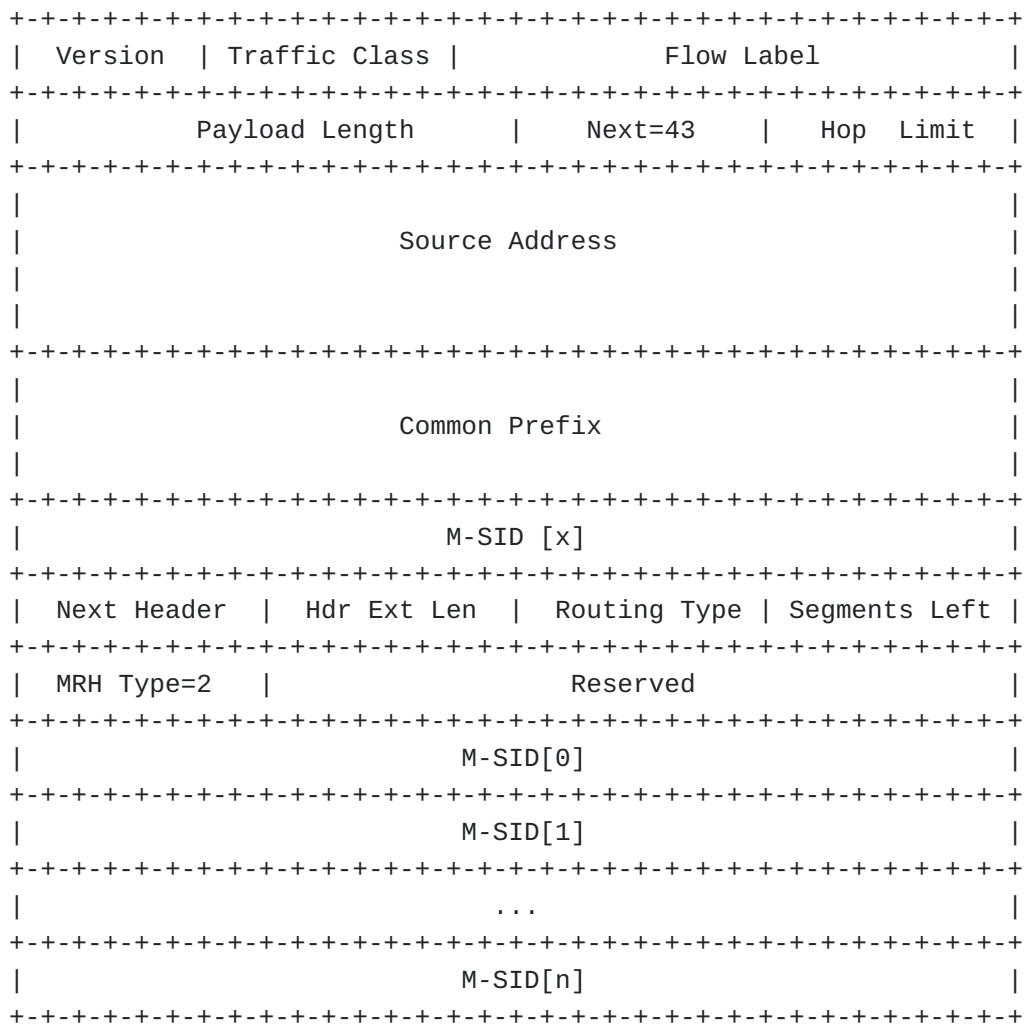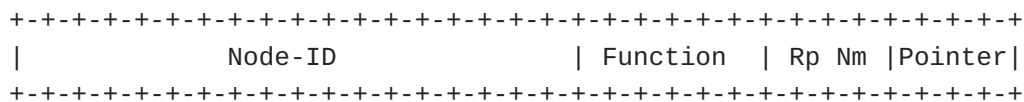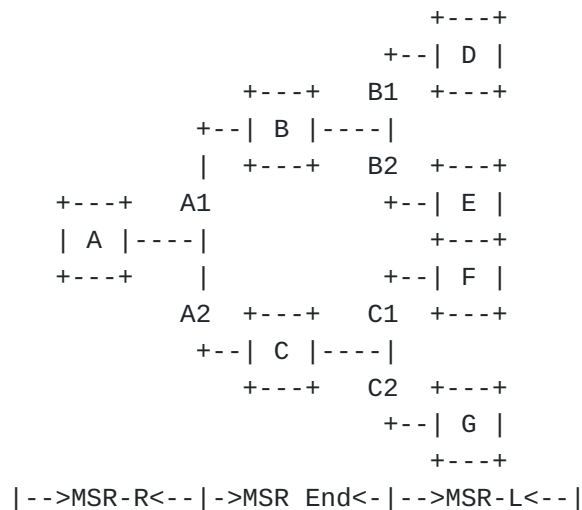points to a Routing Header and the type of the routing header is MRH
Type 2. The Replication Endpoint is supposed to

  *Find the next groups of segments in the MRH which present the
   child nodes of the existing Replication Endpoint with Pointer-1
   and Pointer-2;

  *Replicate the packet based on the number of child nodes, which is
   the replication number;

  *Replace the SID part of IPv6 destination address with the
   corresponding child node's segment in order for each replicated
   packet;

  *Update the segment left field of each replicated packet based on
   the location child node's segment;

  *Forward the packet to the output port based on the FIB with the
   existing destination address as an entry;

8. **Illustration**

   In order to

```
                              +---+
                          +--| D |
                +---+   B1   +---+
            +--| B |----|
            |   +---+   B2   +---+
  +---+   A1             +--| E |
  | A |----|                 +---+
  +---+    |             +--| F |
         A2   +---+   C1   +---+
            +--| C |----|
               +---+   C2   +---+
                          +--| G |
                              +---+
|-->MSR-R<--|->MSR End<-|-->MSR-L<--|
```

   MSR-R: MSR6 Root Node, including Node A;

   MSR End: MSR6 Replication Endpoint, including Node B and C;

   MSR-L: MSR6 Leaf Nodes, including Node D, E, F, G;

8.1. MSR6 Illustration: End.RL

In node A, the packet is encapsulated the packet with an IPv6 header
carrying an MRH. The segment list in the MRH is as follows. The
destination address of the IPv6 header is the 1st SID in the segment
list, which is the local SID of node A. Based on the End.RL behavior
defined in section 5.1, the packet is replicated to 2 (Replication
Number+1=1+1=2) packets. In the 1st packet, the Segment Left is set
to 2 and DA is replaced by Segment List[2]; In the 2nd packet, the
Segment Left is set to 4 and DA is replaced by Segment List[4]; The
packets are routed to the node B and node C respectively.

```
+---------------------------------------+
|   Loc:A | Fun:End.RL | Rp-Nm:1 |  P:2  |
+---------------------------------------+
|   Loc:B | Fun:End.RL | Rp-Nm:1 |  P:4  |
+---------------------------------------+
|   Loc:C | Fun:End.RL | Rp-Nm:1 |  P:6  |
+---------------------------------------+
|   Loc:D | Fun:End.RL | Rp-Nm:0 |  P:0  |
+---------------------------------------+
|   Loc:E | Fun:End.RL | Rp-Nm:0 |  P:0  |
+---------------------------------------+
|   Loc:F | Fun:End.RL | Rp-Nm:0 |  P:0  |
+---------------------------------------+
|   Loc:G | Fun:End.RL | Rp-Nm:0 |  P:0  |
+---------------------------------------+
```

In MSR6 Replication Endpoint B, the destination address of the IPv6
header is the local SID of node B. Based on the End.RL behavior
defined in section 5.1, the packet is replicated to 2 (Replication
Number+1=1+1=2) packets. In the 1st packet, the Segment Left is set
to 4 and DA is replaced by Segment List[4]; In the 2nd packet, the
Segment Left is set to 5 and DA is replaced by Segment List[5]; The
packets are routed to the node D and node E respectively.

In MSR Endpoint C, the process is similar.

In MSRE D, the destination address of the IPv6 header is the local
SID of node D. Based on the End.RL behavior defined in section 5.1,
when replication number=0, node D stops processing the MRH and
continues to process the next header in the packet.

8.2. MSR6 Illustration: End.RL.X

In node A, the packet is encapsulated the packet with an IPv6 header
carrying an MRH. The segment list in the MRH is as follows. The
destination address of the IPv6 header is the local SID of node A.
Based on the End.RL.X behavior defined in section 5.2, the packet is

replicated to 2 (Replication Number+1=1+1=2) packets. In the 1st
packet, the Segment Left is set to 3 and DA is replaced by Segment
List[3] and send the packet based on the specified adjacency A1; In
the 2nd packet, the Segment Left is set to 2 and DA is replaced by
Segment List[2]; Based on the arguments in the updated DA, the
Segment Left is set to 5 and DA is replaced by Segment List[5] and
send the packet based on the specified adjacency A2;

```
+-----------------------------------------+
|  Loc:A1 | Fun:End.RL.X| Rp-Nm:1 |  P:3  |
+-----------------------------------------+
|  Loc:A2 | Fun:End.RL.X| Rp-Nm:0 |  P:5  |
+-----------------------------------------+
|  Loc:B1 | Fun:End.RL.X| Rp-Nm:1 |  P:0  |
+-----------------------------------------+
|  Loc:B2 | Fun:End.RL.X| Rp-Nm:0 |  P:0  |
+-----------------------------------------+
|  Loc:C1 | Fun:End.RL.X| Rp-Nm:1 |  P:0  |
+-----------------------------------------+
|  Loc:C2 | Fun:End.RL.X| Rp-Nm:0 |  P:0  |
+-----------------------------------------+
```

In node B, the packet is encapsulated the packet with an IPv6 header
carrying an MRH. The segment list in the MRH is as follows. The
destination address of the IPv6 header is the local SID of node B.
Based on the End.RL.X behavior defined in section 5.2, the packet is
replicated to 2 (Replication Number+1=1+1=2) packets. In the 1st
packet, the DA is replaced by the corresponding leaf stored in the
node and send the packet based on the specified adjacency B1; In the
2nd packet, the DA is replaced by the corresponding leaf stored in
the node and send the packet based on the specified adjacency B2;

In MSR Endpoint C, the process is similar.

In MSRE D, the destination address of the IPv6 header is the local
SID of node D. Node D stops processing the MRH, and begin to process
the next header in the packet.

9. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an
RFC.

10. Security Considerations

11. Acknowledgements

12. Normative References

[I-D.cheng-spring-ipv6-msr-design-consideration]
          Cheng, W., Mishra, G., Li, Z., Wang, A., Qin, Z., and C.
          Fan, "Design Consideration of IPv6 Multicast Source
          Routing (MSR6)", Work in Progress, Internet-Draft, draft-
          cheng-spring-ipv6-msr-design-consideration-00, 12 July
          2021, <https://www.ietf.org/archive/id/draft-cheng-
          spring-ipv6-msr-design-consideration-00.txt>.

[I-D.geng-bier-ipv6-inter-domain] Geng, L., Xie, J., McBride, M.,
          Yan, G., and X. Geng, "Inter-Domain Multicast Deployment
          using BIERv6", Work in Progress, Internet-Draft, draft-
          geng-bier-ipv6-inter-domain-02, 27 October 2020,
          <https://www.ietf.org/archive/id/draft-geng-bier-ipv6-
          inter-domain-02.txt>.

[I-D.ietf-bier-ping] Kumar, N., Pignataro, C., Akiya, N., Zheng, L.,
          Chen, M., and G. Mirsky, "BIER Ping and Trace", Work in
          Progress, Internet-Draft, draft-ietf-bier-ping-07, 11 May
          2020, <https://www.ietf.org/archive/id/draft-ietf-bier-
          ping-07.txt>.

[I-D.xie-bier-ipv6-mvpn] Xie, J., McBride, M., Dhanaraj, S., Geng,
          L., and G. Mishra, "Use of BIER IPv6 Encapsulation
          (BIERv6) for Multicast VPN in IPv6 networks", Work in
          Progress, Internet-Draft, draft-xie-bier-ipv6-mvpn-03, 10
          October 2020, <https://www.ietf.org/archive/id/draft-xie-
          bier-ipv6-mvpn-03.txt>.

[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>.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the
          Internet Protocol", RFC 4301, DOI 10.17487/RFC4301,
          December 2005, <https://www.rfc-editor.org/info/rfc4301>.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI
          10.17487/RFC4302, December 2005, <https://www.rfc-
          editor.org/info/rfc4302>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC
          4303, DOI 10.17487/RFC4303, December 2005, <https://
          www.rfc-editor.org/info/rfc4303>.

[RFC5374] Weis, B., Gross, G., and D. Ignjatic, "Multicast
          Extensions to the Security Architecture for the Internet

                    Protocol", RFC 5374, DOI 10.17487/RFC5374, November 2008,
                    <https://www.rfc-editor.org/info/rfc5374>.

   [RFC7296]   Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T.
                Kivinen, "Internet Key Exchange Protocol Version 2
                (IKEv2)", STD 79, RFC 7296, DOI 10.17487/RFC7296, October
                2014, <https://www.rfc-editor.org/info/rfc7296>.

   [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>.

   [RFC8279]   Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A.,
                Przygienda, T., and S. Aldrin, "Multicast Using Bit Index
                Explicit Replication (BIER)", RFC 8279, DOI 10.17487/
                RFC8279, November 2017, <https://www.rfc-editor.org/info/
                rfc8279>.

   [RFC8296]   Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A.,
                Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation
                for Bit Index Explicit Replication (BIER) in MPLS and
                Non-MPLS Networks", RFC 8296, DOI 10.17487/RFC8296,
                January 2018, <https://www.rfc-editor.org/info/rfc8296>.

   [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

   Xuesong Geng
   Huawei Technologies

   Email: gengxuesong@huawei.com

   Zhenbin Li
   Huawei Technologies

Email: lizhenbin@huawei.com

Jingrong Xie
Huawei Technologies

Email: xiejingrong@huawei.com