Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 20 September 2026

Y. Liu
China Mobile
J. Xie
X. Geng
Huawei Technologies
M. Chen
New H3C Technologies

19 March 2026

RGB (Replication through Global Bitstring) Segment for Multicast Source Routing over IPv6
draft-lx-msr6-rgb-segment-06

Abstract

This document introduces the RGB (Replication through Global Bitstring) Segment for Multicast Source Routing over IPv6.

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].

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 20 September 2026.

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]. Like SRv6 for unicast, MSR6 provides network programming capability for multicast service by encoding network instructions in the IPv6 packet header, and specifies a packet to replicate and forward based on these instructions. Each instruction is implemented on one or several nodes in the network and identified by an MSR6 Segment Identifier. Similar as BIER, bitstring is used in MSR6 to represent nodes or links in the network to save encapsulation expense.

MSR6 has two basic modes of forwarding: one is based on Shortest Path First(SPF), which is called MSR6 BE(Best Effort) mode; the other is based on traffic engineered, which is called MSR6 TE(Traffic Engineering) mode. This document defines a new type of segment, Replication through Global Bitstring Segment (RGB Segment), and the corresponding packet processing procedures over the IPv6 data plane for the MSR6 BE solutions.

2. Terminologies

The following new terms are used throughout this document:

MSR6 Domain: a set of nodes participating in the multicast source routing over IPv6;

MSR6 Ingress Node: a node through which a multicast data packet enters an MSR6 domain; The MSR6 Ingress Node could be a host or a network device.

MSR6 Egress Node: a node through which a multicast data packet leaves an MSR6 domain; The MSR6 Egress Node could be a host or a network device.

MSR6 Root Node: a node which is the beginning point of a multicast tree for multicast service using MSR6. It encapsulates the packet with an MSR6 multicast header.

MSR6 Leaf Node: a node which is the ending point of a multicast tree for multicast service using MSR6. It decapsulates the MSR6 multicast header in the packet.

MSR6 Replication Endpoint: the intermediate node of a multicast treefor multicast service using MSR6, 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 informated encapsulation in the MSR6 header of the packet.

MSR6 Transit Node: a node which forwards the MSR6 packet as an IPv6 unicast packet between MSR6 replication endpoints or MSR6 replication endpoint and MSR6 leaf node;

3. RGB Destination Options Header

Multicast BE flow, relative to the concept of multicast TE flow, does not need to go through a planning path satisfying service requirements. The path for multicast BE flow is normally provided by routing underlay protocol, as IGP.

In MSR6, a set of egress nodes which the packet is supposed to be sent to are supposed to be indicated in the packet, in order to avoid maintaining multicast tree for each multicast flow. Global bitstring could represent the egress nodes efficiently and such mechanism has been defined in BIER [RFC8279] through BIER Header [RFC8279], which could be reused in MSR6.

For IPv6 data plane, a new IPv6 Destination Options Header ([RFC8200]) option called RGB (Replication through Global Bitstring) is introduced in MSR6.

The encoding of RGB Option is showed as follows:

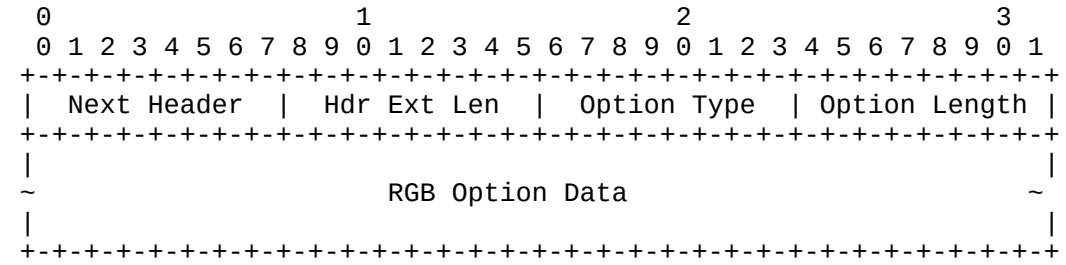

Next Header 8-bit selector. Identifies the type of header immediately following the Destination Options header.

Hdr Ext Len 8-bit unsigned integer. Length of the Destination Options header in 8-octet units, not including the first 8 octets.

Option Type To be allocated by IANA. See section 6.

Option Length 8-bit unsigned integer. Length of the option, in octets, excluding the Option Type and Option Length fields.

The encoding of RGB Option is defined as follows:

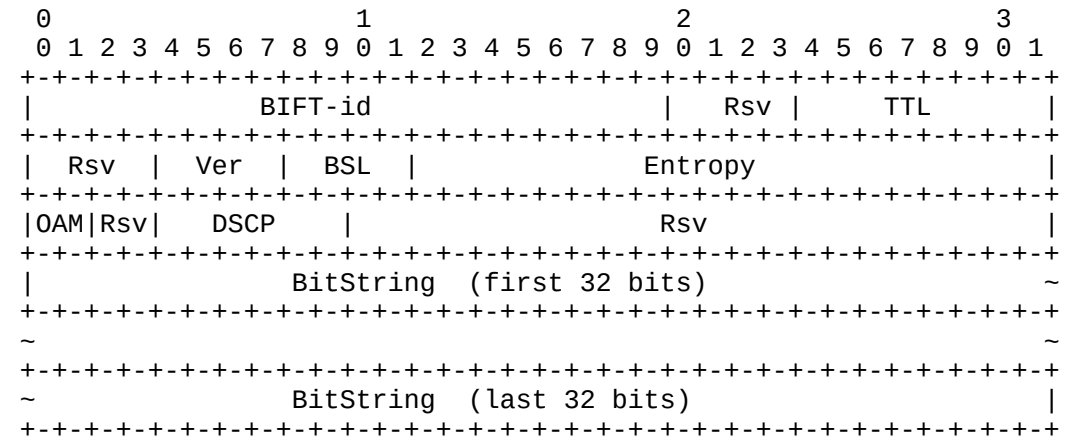

The RGB Option Data reuses some codepoint of Non-MPLS BIER Header defined in [RFC8296] except the fields of Nibble, DSCP and Proto, which are replaced as the Reserved field. The Reserved fields SHOULD be set to 0 and MUST be ignored up reception.

4. RGB Segment

4.1. RGB Segment Definition

As defined in [RFC8402], segment represents instruction, topological or service based.In an IPv6 domain, a segment could be encoded as an IPv6 address.

In MSR6, a new type of segment is defined called RGB segment. RGB segment is used to identify the Replication Endpoint and direct to replicatereplicate the packet using BIER forwarding mechanism defined in [RFC8279] according to the bitstring defined in the RGB Option. RGB segment follows the SID format defined in [RFC8986], consisting of LOC:FUNCT:ARG. RGB segment is advertised by the RGB replication endpoint. In an MSR6 domain, RGB segment is used as the destination address of the MSR BE packet, steering the packet to the next Replication Endpoint. If there is 1 or more MSR6 transit nodes between two Replication Endpoints, the packet is forwarded as normal unicast IPv6 packet, and RGB segment is treated as unicast IPv6 prefix to route the packet

The segment defined in [RFC8402] can represent instruction, topological or service based. In an IPv6 domain, a segment could be encoded as an IPv6 address.

In the IPv6 data plane, RGB segment is a new type of segment which is used to identify the Replication Endpoint. Replication Endpoint is able to replicate the packet using BIER forwarding mechanism according to the bitstring defined in the RGB Option.

RGB segment is used as an IPv6 address, which is 128 bits and follows the SID format defined in [RFC8986], consisting of LOC:FUNCT:ARG. RGB segment is advertised by the RGB replication endpoint.

In an MSR6 domain, RGB segment is used as the destination address of the MSR BE packet, when a packet is replicated to the next Replication Endpoint. If there is 1 or more MSR6 transit node between two Replication Endpoints, the packet is forwarded as normal unicast IPv6 packet.

4.2. End.RGB Behavior

In SRv6, a packet processing behavior is executed at the SRv6 Segment Endpoint Node ([RFC8986]). Similarly, in MSR6, a new type of behavior, End.RGB(End. Replication through Global Bitstring), is defined for RGB Segment. The pseudo-code for End.RGB is defined in this section.

When an MSR6 Replication Endpoint receives a packet whose IPv6 DA (Destination Address) is a SID and the SID is a local End.RGB SID, the MSR6 Replication Endpoint does the following:

```
1. IF (There is DoH as an IPv6 Extension header and one of the options type is RGB);
2.   Lookup BIFT(Bit Index Forwarding Table, RFC8279) based on the bitstring
       inside the RGB Option Data.
3.   Forward the packet via the matched entry in the BIFT.
4. ELSE IF NH=ICMPv6 or (NH=RGB Extension Header Type and NH
    of Extension Header=ICMPv6)  ;
5.   Send to CPU.
6. ELSE  ;Ref
7.   Drop the packet.
```

Ref: An ICMPv6 packet using End.RGB as destination address.

5. MSR6 BE Encapsulation

MSR6 BE encapsulation is composed of 3 parts: IPv6 header, IPv6 RGB option DoH and client multicast packet, which is showedas follows:

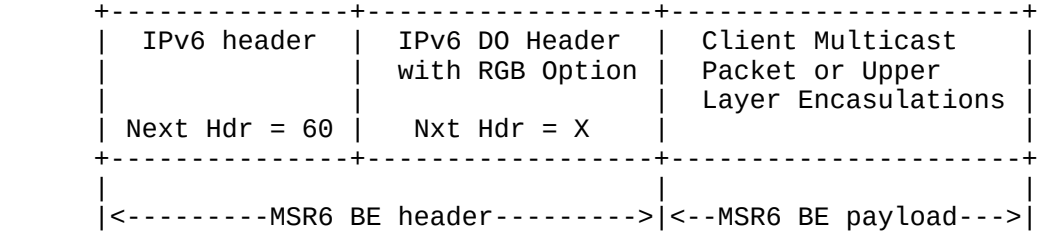

In the MSR6 BE header, the RGB Segment is used as the IPv6 Destination Address and indicates the next MSR6 Replication Endpoints in an MSR domain. RGB DoH option is used as the carrier of bitstring information and the MSR6 Replication Endpoint uses the bitstring as the entry to look up BIFT(Bit Index Forwarding Table) to replicate and find the next MSR6 Replication Endpoints.

6. Packet Processing Procedure

This section defines the general process of MSR6 BE to transport a multicast service. The corresponding control plane is out of scope of this document and could be discussed in the following work.

MSR6 Root Node: The Ingress Node of a multicast flow. It can be either a host originating the MSR6 packet, or a router encapsulating the customer packet in an MSR6 header. The bitstring in the DoH is determined by the egress nodes the packet is supposed to be replicated to. The IPv6 destination address is the RGB segment which is determined by the next MSR6 RGB Replication Endpoints the packet

is supposed to be sent to. The downstream MSR6 Replication Endpoints are determined by the matched entries in BIFT according to the bier forwarding mechanism.

MSR6 Replication Endpoint: Replicate the packet and forward the packet to the next MSR6 Replication Endpoints. When an MSR6 Replication Endpoint receives a packet whose IPv6 Destination Address is A and A is the local RGB SID for the existing MSR6 Replication Endpoint, process the bitstring in the RGB DoH of the packet and look up the corresponding BITF for the next MSR6 Replication Endpoints. Replicate the packet, update the bitstring and DA in each replicated packet based on the lookup result.The RGB processing procedure follows the specification in BIER architecture defined in ' [RFC8279].

MSR6 Transit Node: Transit the packet as a unicast IPv6 packet by looking up FIB until find the next MSR6 Replication Endpoint.

MSR6 Leaf Node: The Egress Node of a multicast flow. When an MSR6 Replication Endpoint receives a packet whose IPv6 Destination Address is A and A is the local RGB segment and the one of the bits which is set to 1 identifies the MSR6 the egress node. If the MSR6 egress node is the edge of a network domain, copy the packet and send the copy to the multicast flow overlay; If the MSR6 egress node is the host supposed to receive the packet, send the packet to the upper layer.

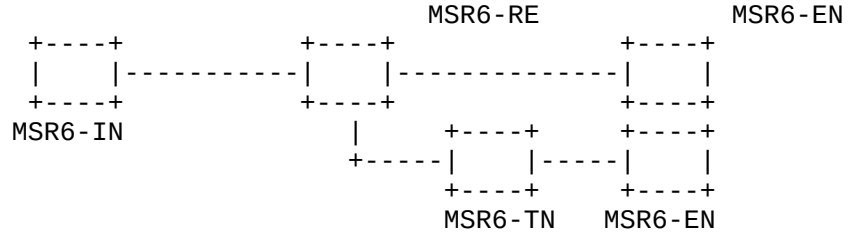

MSR6-IN: MSR6 Ingress Node (MSR6 Root Node)
MSR6-TN: MSR6 Transit Node (which is not MSR6 aware)
MSR6-RE: MSR6 Replication Endpoint
MSR6-EN: MSR6 Egress Node (MSR6 Leaf Node)

7. Illustration

* Case 1: Host originating MSR6 BE

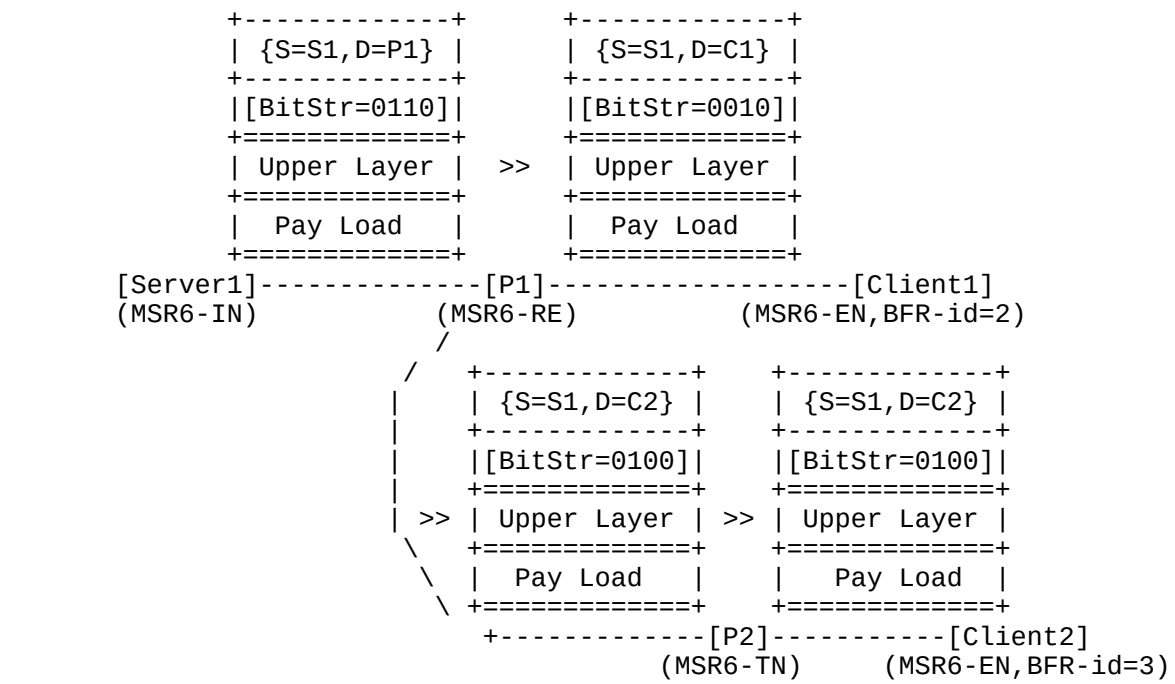

{S=Server1,D=P1}: Source address and Destination address in IPv6 header.
[BitStr=0110]: BitString value in IPv6 Destination Options Header.
MSR6-IN: MSR6 Ingress Node
MSR6-TN: MSR6 Transit Node (which is not MSR6 aware)
MSR6-RE: MSR6 Replication Endpoint
MSR6-EN: MSR6 Egress Node

Server1 generates the packet with an IPv6 Header. Knowing that BFR-ID of Client 1 is 2 and BFR-ID of Client 2 is 3, it follows that when the multicast service is supposed to be transmitted to Client1 and Client2, the bitstring in RGB DoH of the IPv6 header is set as "0110". Look up the BIFT and finds the RGB segment of next MSR6 BFR is P1. The IPv6 DA is set as "P1".

P1 receives the packet with DA as "P1", which is the local RGB segment. P1 parses the DoH with RGB Option Data and looks up the BIFT to find the corresponding entry. P1 replicates the packets into 2 copies based on the look up result. DA of one replicated packet is set to "C1" and the bitstring is set to "0010". DA of the other replicated packet is set to "C2" and the bitstring is set to "0100". These 2 packets are forwarded to next hop based on the updated DA.

P2 receives the packet and forwards it Client2 based on the DA of "C2".

Client1 receives the packet with DA as "C1". "C1" is the local RGB segment and "0010" identifies Client1 itself. The packet is sent to the upper layer.

Client2 receives the packet with DA as "C2". "C2" is the local RGB segment and "0100" identifies Client2 itself. The packet is sent to the upper layer.

* Caes 2: MSR6 is used in a network domain

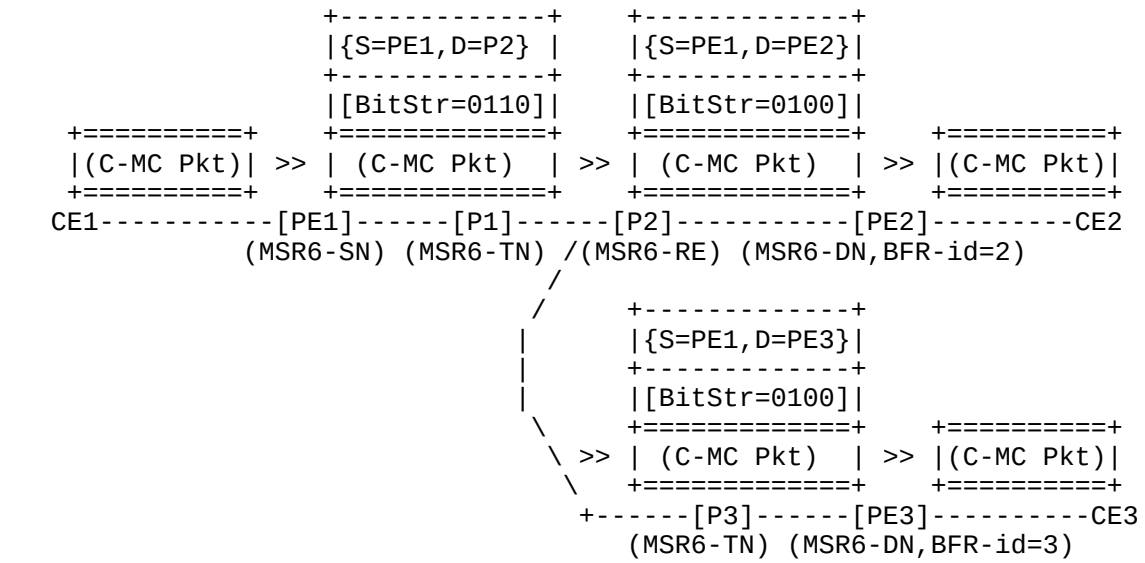

{S=CE1,D=P1}: Source address and Destination address in IPv6 header.
[BitStr=0110]: BitString value in IPv6 Destination Options Header.
(C-MC Pkt): Customer MultiCast packet.
MSR6-IN: MSR6 Ingress Node
MSR6-TN: MSR6 Transit Node (which is not MSR6 aware)
MSR6-RE: MSR6 Replication Endpoint
MSR6-EN: MSR6 Egress Node

PE1 receives the customer multicast packet from CE1. An MSR BE header is encapsulated as defined in section 3. Knowing that BFR-ID of PE 1 is 2 and BFR-ID of PE 2 is 3, it follows that when the multicast service is supposed to be transmitted to PE2 and PE3, the bitstring in the RGB Options Header of DoH is set as "0110". Look up the corresponding BIFT and finds the RGB segment of next MSR6 BFR is P2. The IPv6 DA is set as "P2".

P1 receives the packet and forwards it P2 based on the DA of "P2".

P2 receives the packet with DA as "P2", which is the local RGB segment. P2 parses the DoH with RGB Option Data and looks up the BIFT to find the corresponding entry.P2 replicates the packets into 2 copies based on the look up result. DA of one replicated packet is set to "PE2" and the bitstring is set to "0010". DA of the other replicated packet is set to "PE3" and the bitstring is set to "0100". These 2 packets are forwarded to next hop based on the updated DA.

P3 receives the packet and forwards it PE3 based on the DA of "PE3".

PE2 receives the packet with DA as "PE2". "PE2" is the local RGB segment and "0010" identifies PE2 itself. The packet is sent to the multicast flow overlay.

PE3 receives the packet with DA as "PE3". "PE3" is the local RGB segment and "0100" identifies PE3 itself. The packet is sent to the multicast flow overlay.

8. IANA Considerations

8.1. RGB Option Type

Allocation is expected from IANA for a RGB Option Type codepoint from the "Destination Options and Hop-by-Hop Options" sub-registry of the "Internet Protocol Version 6 (IPv6) Parameters" registry.

Hex Value	act	chg	rest	Description	Reference
TBD	01	1	TBD	RGB Option	This draft

8.2. End.RGB Function

Allocation is expected from IANA for an End.RGB function codepoint from the "SRv6 Endpoint Behaviors" sub-registry. The value 60 is suggested.

Value	Hex	Endpoint function	Reference
TBD	TBD	End.RGB	This draft

9. Security Considerations

The MSR6 domain can be a single IGP area, an anonymous system (AS) with multiple IGP areas, or multiple anonymous systems (ASes) operated by a network operator.

It is expected that all nodes in an MSR6 domain are managed by the same administrative entity. MSR6-encapsulated packets should generally not be accepted from untrusted interfaces or tunnels. For example, an operator may wish to have a policy of accepting MSR6 encapsulated packets only from interfaces to trusted routers, and not from customer-facing interfaces.

For applications that require a MSR6 Replication Endpoint to accept a MSR6 encapsulated packet from an interface to a system that is not controlled by the network operator, the security considerations of [RFC8296] apply

9.1. Intra Domain Deployment

Generally nodes outside the MSR6 Domain are not trusted: they cannot directly use the End.RGB segment of the domain. This is enforced by two levels of access control lists:

1. Any packet entering the MSR6 Domain and destined to an End.RGB Segment within the MSR6 Domain is dropped. This may be realized with the following logic. Other methods with equivalent outcome are considered compliant:

* allocate all the End.RGB Segment from a block S/s

* configure each external interface of each edge node of the domain with an inbound infrastructure access list (IACL) which drops any incoming packet with a destination address in S/s

* Failure to implement this method of ingress filtering may expose the MSR6 Domain to BIER attacks. The security consideration on BIER attacks is as described and referenced in [RFC8296].

2. The distributed protection in #1 is complemented with per node protection, dropping packets to End.RGB Segment from source addresses outside the MSR6 Domain. This may be realized with the following logic. Other methods with equivalent outcome are considered compliant:

* assign all interface addresses from prefix A/a

* assign all the IPv6 addresses used as source address of MSR6 packets from a block B/b

* at node k, all End.RGB Segment IPv6 addresses local to k are assigned from prefix Sk/sk

* configure each internal interface of each MSR6 node k in the MSR6 Domain with an inbound IACL which drops any incoming packet with a destination address in Sk/sk if the source address is not in A/a or B/b.

For simplicity of deployment, a configuration of IACL effective for all interfaces can be provided by a router. Such IACL can be referred to as global IACL(GIACL) .Each MSR6 node k then simply configures a GIACL which drops any incoming packet with a destination address in Sk/sk if the source address is not in A/a or B/b for the intra-domain deployment mode.

9.2. ICMP Error Processing

The MSR6 Replication Endpoint does not send ICMP error messages to the source address of a MSR BE packet, but there is still chance that Non-MSR6 Replication Endpoint routers send ICMP error messages to source nodes within the MSR6 Domain.

A large number of ICMP may be elicited and sent to a MSR6 Ingress router, in case when an MSR6 BE packet is filled with wrong Hop Limit, either error or malfeasance. A rate-limiting of ICMP packet should be implemented on each MSR6 Replication Endpoint.

The ingress node can take note of the fact that it is getting, in response to MSR6 BE packet, one or more ICMP error packets. By default, the reception of such packet MUST be countered and logged. However, it is possible for such log entries to be "false positives" that generate a lot of "noise" in the log; therefore, implementations SHOULD have a knob to disable this logging.

9.3. Security caused by RGB option

This document introduces a new option used in IPv6 Destination Options Header. An IPv6 packet with a normal IPv6 address of a router (e.g. loopback IPv6 address of the router) as destination address will possibly carry a RGB option.

For a router incapable of MSR6 BE, such MSR6 BE packet will not be processed by the procedure described in this document, but be processed as normal IPv6 packet with unknown option, and the existing security considerations for handling IPv6 options apply. Possible way of handling IPv6 packets with RGB option may be send to CPU for slow path processing, with rate-limiting, or be discarded according to the local policy.

For a router capable of MSR6 BE, such MSR6 BE packet MUST NOT be forwarded, but should be processed as a normal IPv6 packet with unknown option, or additionally and optionally be countered and logged if the router is capable of doing so.

10. Normative References

[I-D.cheng-spring-ipv6-msr-design-consideration]
Cheng, W., Mishra, G. S., 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-01, 25 October 2021, <https://datatracker.ietf.org/doc/html/draft-cheng-spring-ipv6-msr-design-consideration-01>.

[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>.

[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>.

[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

Yisong Liu
China Mobile
Email: liuyisong@chinamobile.com

Jingrong Xie
Huawei Technologies
Email: xiejingrong@huawei.com

Xuesong Geng
Huawei Technologies
Email: gengxuesong@huawei.com

Mengxiao Chen
New H3C Technologies
Email: chen.mengxiao@h3c.com