Workgroup: Network Working Group
Internet-Draft:
draft-anish-spring-bimap-multicast-00
Published: 26 February 2023
Intended Status: Standards Track
Expires: 30 August 2023
Authors: A. Peter, Ed.
Individual Contribution

SRv6 Bitmap Multicast.

Abstract

Multicast forwarding in a network provides advantages in improving the network usage and performance. In some cases it helps improve the operations in managing network. The major challenge in multicast operations is in managing the per-flow states in the network as required by all the legacy multicast frameworks.

This document specifies a bitmap forwarding extension to SRv6 to support state-free forwarding model in a network.

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 RFC2119 [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 30 August 2023.

Copyright Notice

Table of Contents

1. Introduction

Segment routing with IPv6 as specified in RFC8689 [RFC8689] provides a source-routing solution for next generation network requirements. More applications and use-cases are finding a better solutions using SRv6. Along with this comes the need to support multicasting and broadcasting in such networks. The various use-cases for this would be stated in the subsequent sections.

Broadcasting typically needs a point-to-multipoint (p2mp) distribution with all the nodes in the network being receivers. Multicasting would imply p2mp distribution along with multipoint-to-multipoint (mp2mp) packet distribution with the participants being pre-determined via a discovery or provisioning mechanism.

Bit-Index-Explicit-Replication specified in RFC8279 [RFC8279] introduced a per-flow-state-free forwarding for multicast using a bit-indexed addressing of multicast receivers.

This document introduces a bit-map based distribution schema in IPv6 networks to achieve p2mp distribution patterns. SRv6 introduced a new semantic to IPv6 address by fragmenting the address space into Locator:Function:Argument construct to achieve the desired SR functionality. This document proposes a similar treatment of IPv6 address to achieve BIER forwarding.

2. IPv6 Bit-Index Format

This document provides a new semantic to the IPv6 address as SI_LOCATOR:BITSTRING:FUNCTION:ARGUMENT. This structure is partly borrowed from SRv6. The BITSTRING part is newly introduced to address the egress routers in the BIER subdomain by its bit index. From here on this format is called as Bit-Index-6 format (BI6)

BIER architecture envisages forwarding by identifying each egress router with an BFR-id. These BFR-id in forwarding translates to a Set-Identifier (SI) and Bitstring. In the IPv6 Bit-Index format, the SI is identified by the SI_LOCATOR and bitstring is encoded in the BITSTRING part of the BI6. The FUNCTION and ARGUMENT bits are part of the format. But depending on the network requirement their lengths may be set to 0 for using this bits for extended bitstring.

SI_LOCATOR is defined as a routable prefix to reach the specific set of routers in a SetIdentifier. Once a BI6 packet reaches a router that is part of a SI, The bit-index based part is referred to for forwarding towards the BFER's with the BIER forwarding principles. The semantics of the FUNC and ARG bits is global in the Sub-domain. The attributes of FUNCTION and ARGUMENT bits must be pre-determined.

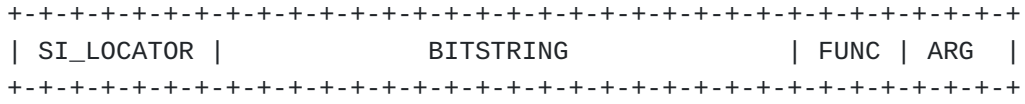

Figure 1: Syntax of BIER6 address

3. Network Overview

BIER architecture puts forward a multicast forwarding based on "Bit-Index-Explicit-Replication". This architecture defines a BIER domain in which an ingress router would encapsulate p2mp packet with a BIER header RFC8296 [RFC8296] . This BIER packet would be replicated to the egress routers identified by the ingress in its BIER header,

over an optimal per-flow-stateless tree discovered with the underlay.

4. Use-Cases

5. Subscriber management

In BIER architecture the multicast egress routers must be learned by the ingress router. This discovery may happen via some out-of-band mechanism beyond the scope of this document.

6. Interworking with non compatible BI6 Routers

A network topology may have legacy devices which may not be capable of BI6 processing. When deploying BI6 the traffic may have to pass through some of these devices for loop-free forwarding.

A router may come to know about the BI6 capability of a neighbouring device via the capabilities it has published in its IGP advertisement. Based on this IGP may form a map of BFER to the nearest BFR on the path to egress. If the BFR is not directly connected, then that BFR's node sid may be inserted into the SRH prior to forwarding the packet.

7. Scope for future work

7.1. Routing extension header for BIER

7.2. Define egress functions based on FUNC and ARG bits

7.3. IGP extension to support underlay

7.4. Discovery mechanism for receivers

8. Management Considerations

9. IANA Considerations

This specification introduces new semantics for IPv6 address. Though this draft does not need any allocations, New IANA allocations would be required for the supplimentary specifications.

10. Security Considerations

This document proposes a semantic for IPv6 address. The security challenges that apply to IPv6 and in the BIER architecture applies to the intended BI6 forwarding model specified here.

The further security scenarios would be added in the due course.

11. Appendix 1: Bit-Index string length

11.1. Private IPv6 address for operations

The Unique Local IPv6 address allocation RFC4193 [RFC4193] provides free to use address blocks with SI_LOCATOR size of 48. This provides a maximum BI6 addressing space of 80 bit length.

The Bit-index string length that can be used would be determined by the SI_locator prefix length and the need for FUNC and ARG bits. Hence if Unique local address space is used, upto 80 BFER's can be addressed.

12. Appendix 2: Scaling considerations

With this specification the typical scale a BIER domain would be less than 96 egress routers. On networks with larger scale the current proposal is to have multiple subdomains and to do ingress replication for traffic bound to various subdomains.

13. References

13.1. 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>.

[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>.

13.2. Informative References

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <https://www.rfc-editor.org/info/rfc4193>.

[RFC8689] Fenton, J., "SMTP Require TLS Option", RFC 8689, DOI 10.17487/RFC8689, November 2019, <https://www.rfc-editor.org/info/rfc8689>.

Author's Address

Anish Peter (editor)
Individual Contribution
Bangalore 560043
KA
India

Email: anish.ietf@gmail.com