

**An IPv6 based BIER Routing Underlay
draft-pfister-bier-over-ipv6-00**

Abstract

This document specifies a routing underlay which enables transport of IPv6 payloads to multiple IPv6 destinations using the Bit Index Explicit Replication (BIER). The BIER bits are stored within the low-order bits of the IPv6 destination address while the high-order bits are left untouched and used in order to decide whether an IPv6 packet is a regular IPv6 packet or an IPv6 BIER packet, and if so, to identify the BIER set identifier.

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 <http://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 March 6, 2017.

Copyright Notice

Copyright (c) 2016 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 (<http://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 Simplified BSD License text as described in [Section 4](#).e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	2
3.	IPv6 BIER Packet Format	3
4.	BIER Layer Operations	3
5.	Routing Underlay Operations	4
6.	Advantages of this Technique	4
7.	Security Considerations	5
8.	IANA Considerations	5
9.	References	5
9.1.	Normative References	5
9.2.	Informative References	5
Appendix A.	Acknowledgements	5
	Authors' Addresses	5

[1.](#) Introduction

The Bit Index Explicit Replication (BIER - [\[I-D.ietf-bier-architecture\]](#)) forwarding technique enables IP multicast transport across a BIER domain. It operates using three different elements, a multicast flow overlay, a BIER Layer, and a routing underlay. This document specifies a way to send IPv6 payloads to multiple destinations using regular IPv6 packets with no additional extension header.

BIER bits are encoded in the low-order bits of the IPv6 destination address of each packet. The high-order bits of the IPv6 destination address are left untouched and used by intermediate BIER routers to decide whether the packet should be forwarded as a regular IPv6 packet or an IPv6 BIER packet, and if so, to know the BIER Set Identifier.

Transported payloads can be of many types such as IPv6 or IPv4 unicast or multicast packets (e.g. using generic packet tunneling [\[RFC2473\]](#)), or transported data (e.g. using UDP). Special care must be taken when forwarding some types of payloads. For example, the UDP checksum must be recomputed when the BIER bits are changed.

[2.](#) Terminology

In this document, the key words "MAY", "MUST", "MUST NOT", "RECOMMENDED", and "SHOULD", are to be interpreted as described in [\[RFC2119\]](#).

3. IPv6 BIER Packet Format

BIER IPv6 packets are IPv6 unicast packets. There is no specificities to the format of the packet. The BIER bits are encoded in the low-order bits of the IPv6 destination address of the packet while the high-order bits are used by intermediate routers to identify that the forwarded packet is an IPv6 BIER packet, and which BIER Set Identifier to use.




Figure 1: IPv6 BIER destination address format

BIER IPv6 Prefix: This is a prefix used for BIER forwarding within the domain. BIER routers will consider all packets sent to this prefix as IPv6 BIER packets.

SI bits: These bits are used to encode the Set Identifier. It is used in order to address more BFERS than there are bits in the last part of the IPv6 address.

BitString bits: Set of bits encoding the set of BFERS the packet should be sent to. Those bits are modified as the packet is replicated by intermediate BIER routers. The length of this field is equal to 128-p-i and specifies the value of the BitStringLength to be used in the BIER forwarding process.

For example, a BIER domain assigns the prefix 2001:db8::/56 as a BIER IPv6 Prefix, and sets i to 8. Every BIER router within the domain is configured with such values and is therefore able to identify all BIER packets and treat them accordingly. As the SI bits length is 8, the lowest-weight 64 bits are used as a BIER bit field. A single packet can be sent to at most 64 different BFERS, while the entire domain can include up to $64 * 2^8 = 16384$ BFERS.

4. BIER Layer Operations

When a multicast packet enters the BIER domain, the BFIR first consults the multicast flow overlay and obtains the set of BFERS the packet must be sent to. This set is used in order to compute the set of bit indexes representing the set of destination BFERS. All indexes that have the same SI are grouped in order to create a set of BitString associated with their respective SI. For each SI, the multicast packet is encapsulated within an IPv6 BIER packet, as specified in [Section 3](#).

The same process is used when a given IPv6 payload should be sent to a set of destinations. But instead of encapsulating the packet, the payload is attached to the BIER IPv6 header and the IPv6 protocol number is set to the type of the payload.

5. Routing Underlay Operations

A BIER router is configured with at least one, and possibly more, rules. Each rule includes a BIER IPv6 Prefix (value and prefix length), the value of 'i', and the set of mappings used by BIER in order to associate BIER bits with IPv6 destinations.

For each packet forwarded by a BFR, the router first performs a longest prefix match operation over the destination address. If the found entry corresponds to a BIER rule, the packet is treated as a BIER packet. Otherwise, it is treated as a regular IPv6 packet.

When an IPv6 BIER packet is forwarded, the BFR retrieves the SI value and BitString and performs the BIER forwarding algorithm. For each replicated packet, the BitString is possibly modified and the packet is sent on the outgoing interface.

6. Advantages of this Technique

The technique described in this document offers different advantages:

BIER IPv6 packets are regular IPv6 packets. If the BIER IPv6 Prefix is a globally unique IPv6 prefix, reachable from outside the BIER domain, it is possible to send a packet from outside the BIER domain to multiple destination within the BIER domain.

It may be used for transporting IP multicast packets, but also for sending IP payloads directly to multiple destinations.

It does not rely on a new IPv6 extension header, which simplifies deployment and is likely to improve performances.

It makes use of a typical IP longest match in order to decide whether a packet is a BIER packet or not, which means hardware and software existing solutions may be used for that purpose.

It is possible to configure a host with an address which corresponds to a BIER address with a single bit set. From the host perspective, such address is not different from a regular IPv6 address. Which means a BIER-unaware host may receive BIER packets transparently (With possible additional duplication by the last BFR when multiple hosts are located on the same link).

7. Security Considerations

This technique allows IPv6 BIER packets to be sent across the internet toward multiple destination located in a given BIER domain. If this is considered a threat, a firewall at the entrance of the BIER domain in order to avoid BIER packets from being injected and replicated within the network.

8. IANA Considerations

This specification does not require any action from IANA.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", [RFC 2473](#), DOI 10.17487/RFC2473, December 1998, <<http://www.rfc-editor.org/info/rfc2473>>.

9.2. Informative References

- [I-D.ietf-bier-architecture]
Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", [draft-ietf-bier-architecture-01](#) (work in progress), June 2015.

Appendix A. Acknowledgements

Comments concerning this document are very welcome.

Authors' Addresses

Pierre Pfister
Cisco Systems
Paris
France

Email: pierre.pfister@darou.fr

IJsbrand Wijnands
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium

Email: ice@cisco.com