

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 26, 2019

X. Xu
Alibaba Inc.
G. Shepherd
Cisco
January 22, 2019

Encapsulating Non-MPLS-BIER in UDP
draft-xu-bier-non-mpls-encap-over-udp-04

Abstract

Bit Index Explicit Replication (BIER) is a new multicast forwarding paradigm which doesn't require an explicit tree-building protocol nor intermediate routers to maintain any multicast state. BIER has two types of encapsulation formats: one is MPLS-BIER encapsulation, the other is non-MPLS-BIER encapsulation. This document proposes a mechanism of encapsulating non-MPLS-BIER packets over UDP tunnels.

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 July 26, 2019.

Copyright Notice

Copyright (c) 2019 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 (<https://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
1.1.	Requirements Language	2
2.	Terminology	2
3.	Encapsulation in UDP	3
4.	Processing Procedures	4
5.	Congestion Considerations	5
6.	Applicability Statements	5
7.	Acknowledgements	5
8.	IANA Considerations	5
9.	Security Considerations	6
10.	References	6
10.1.	Normative References	6
10.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

Bit Index Explicit Replication (BIER) [[RFC8279](#)] is a new multicast forwarding paradigm which doesn't require an explicit tree-building protocol nor intermediate routers to maintain any multicast state. As described in [Section 6.9 of \[RFC8279\]](#), a BFR may need to tunnel a BIER packet over a certain kind of tunnel, e.g., UDP tunnel.

[RFC8296] defines two types of BIER encapsulation formats: one is MPLS-BIER encapsulation, the other is non-MPLS-BIER encapsulation. MPLS-BIER packets can be transported over UDP tunnels by using the MPLS-in-UDP encapsulation as described in [[RFC7510](#)]. This document proposes a mechanism of encapsulating non-MPLS-BIER packets over UDP tunnels.

[1.1.](#) 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](#)].

[2.](#) Terminology

This memo makes use of the terms defined in [[RFC8279](#)] and [[RFC8296](#)].

3. Encapsulation in UDP

Non-MPLS-BIER-in-UDP encapsulation format is shown as follows:




Figure 1: Non-MPLS-BIER-in-UDP Encapsulation Format

Source Port of UDP:

This field contains a 16-bit entropy value that is generated by the encapsulator to uniquely identify a flow. What constitutes a flow is locally determined by the encapsulator and therefore is outside the scope of this document. What algorithm is actually used by the encapsulator to generate an entropy value is outside the scope of this document. For example, the 20-bit entropy value contained in the BIER header could actually be transformed to a 16-bit value and then be filled into this field.

In case the tunnel does not need entropy, this field of all packets belonging to a given flow **SHOULD** be set to a randomly selected constant value so as to avoid packet reordering.

To ensure that the source port number is always in the range 49152 to 65535 (Note that those ports less than 49152 are reserved by IANA to identify specific applications/protocols) which may be required in some cases, instead of calculating a 16-bit hash, the encapsulator **SHOULD** calculate a 14-bit hash and use those 14 bits as the least significant bits of the source port field while the most significant two bits **SHOULD** be set to binary 11. That still conveys 14 bits of entropy information which would be enough as well in practice.

Destination Port of UDP:

This field is set to a value (TBD1) allocated by IANA to indicate that the UDP tunnel payload is a non-MPLS-BIER packet.

UDP Length:

The usage of this field is in accordance with the current UDP specification [[RFC0768](#)].

UDP Checksum:

For IPv4 UDP encapsulation, this field is RECOMMENDED to be set to zero for performance or implementation reasons because the IPv4 header includes a checksum and use of the UDP checksum is optional with IPv4. For IPv6 UDP encapsulation, the IPv6 header does not include a checksum, so this field MUST contain a UDP checksum that MUST be used as specified in [[RFC0768](#)] and [[RFC2460](#)] unless one of the exceptions that allows use of UDP zero-checksum mode (as specified in [[RFC6935](#)]) applies.

Non-MPLS-BIER Packet:

This field contains one non-MPLS-BIER packet.

4. Processing Procedures

This Non-MPLS-BIER-in-UDP encapsulation causes non-MPLS BIER packets to be forwarded across an IP transit core via "UDP tunnels". While performing Non-MPLS-BIER-in-UDP encapsulation, an encapsulator would generate an entropy value and encode it in the Source Port field of the UDP header. The Destination Port field is set to a value (TBD1) allocated by IANA to indicate that the UDP tunnel payload is a non-MPLS-BIER packet. Transit routers, upon receiving these UDP encapsulated non-MPLS-BIER packets, could balance these packets based on the hash of the five-tuple of UDP packets. Decapsulators receiving these UDP encapsulated non-MPLS-BIER packets MUST decapsulate these packets by removing the UDP header and then forward them accordingly.

Similar to all other IP-based tunneling technologies, Non-MPLS-BIER-in-UDP encapsulation introduces overheads and reduces the effective Maximum Transmission Unit (MTU) size. Non-MPLS-BIER-in-UDP encapsulation may also impact Time-to-Live (TTL) or Hop Count (HC)

and Differentiated Services (DSCP). Hence, Non-MPLS-BIER-in-UDP MUST follow the corresponding procedures defined in [[RFC2003](#)].

Encapsulators MUST NOT fragment non-MPLS-BIER packet, and when the outer IP header is IPv4, encapsulators MUST set the DF bit in the outer IPv4 header. It is strongly RECOMMENDED that IP transit core be configured to carry an MTU at least large enough to accommodate the added encapsulation headers. Meanwhile, it is strongly RECOMMENDED that Path MTU Discovery [[RFC1191](#)] [[RFC1981](#)] or Packetization Layer Path MTU Discovery (PLPMTUD) [[RFC4821](#)] is used to prevent or minimize fragmentation.

5. Congestion Considerations

As it's explicitly stated in the Application Statements ([Section 6](#)), this Non-MPLS-BIER-in-UDP encapsulation method MUST only be used within networks that are well-managed, therefore, congestion control mechanism is not needed.

6. Applicability Statements

This Non-MPLS-BIER-in-UDP encapsulation technology MUST only be used within networks which are well-managed by a service provider and MUST NOT be used within the Internet. In the well-managed network, traffic is well-managed to avoid congestion and fragmentation on encapsulated packets (i.e., Non-MPLS-BIER packets) are not needed.

7. Acknowledgements

TBD.

8. IANA Considerations

One UDP destination port number indicating non-MPLS-BIER needs to be allocated by IANA:

Service Name: Non-MPLS-BIER-in-UDP Transport Protocol(s):UDP
Assignee: IESG <iesg@ietf.org>
Contact: IETF Chair <chair@ietf.org>.
Description: Encapsulate Non-MPLS-BIER packets in UDP tunnels.
Reference: This document.
Port Number: TBD1 -- To be assigned by IANA.

One UDP destination port number indicating Non-MPLS-BIER with DTLS needs to be allocated by IANA:

Service Name: Non-MPLS-BIER-in-UDP-with-DTLS

Transport Protocol(s): UDP

Assignee: IESG <iesg@ietf.org>

Contact: IETF Chair <chair@ietf.org>.

Description: Encapsulate Non-MPLS-BIER packets in UDP tunnels with DTLS.

Reference: This document.

Port Number: TBD2 -- To be assigned by IANA.

9. Security Considerations

The security problems faced with the Non-MPLS-BIER-in-UDP tunnel are exactly the same as those faced with MPLS-in-UDP tunnel [RFC7510]. In other words, the Non-MPLS-BIER-in-UDP tunnel as defined in this document by itself cannot ensure the integrity and privacy of data packets being transported through the Non-MPLS-BIER-in-UDP tunnel and cannot enable the tunnel decapsulator to authenticate the tunnel encapsulator. In the case where any of the above security issues is concerned, the Non-MPLS-BIER-in-UDP tunnel SHOULD be secured with IPsec or DTLS. IPsec was designed as a network security mechanism and therefore it resides at the network layer. As such, if the tunnel is secured with IPsec, the UDP header would not be visible to intermediate routers anymore in either IPsec tunnel or transport mode. As a result, the meaning of adopting the Non-MPLS-BIER-in-UDP tunnel as an alternative to the Non-MPLS-BIER-in-GRE or Non-MPLS-BIER-in-IP tunnel is lost. By comparison, DTLS is better suited for application security and can better preserve network and transport layer protocol information. Specifically, if DTLS is used, the destination port of the UDP header will be filled with a value (TBD2) indicating non-MPLS-BIER with DTLS and the source port can still be used as an entropy field for load-sharing purposes.

10. References

10.1. Normative References

- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, [RFC 768](#), DOI 10.17487/RFC0768, August 1980, <<https://www.rfc-editor.org/info/rfc768>>.
- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", [RFC 1191](#), DOI 10.17487/RFC1191, November 1990, <<https://www.rfc-editor.org/info/rfc1191>>.
- [RFC1981] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", [RFC 1981](#), DOI 10.17487/RFC1981, August 1996, <<https://www.rfc-editor.org/info/rfc1981>>.

- [RFC2003] Perkins, C., "IP Encapsulation within IP", [RFC 2003](#), DOI 10.17487/RFC2003, October 1996, <<https://www.rfc-editor.org/info/rfc2003>>.
- [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>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), DOI 10.17487/RFC2460, December 1998, <<https://www.rfc-editor.org/info/rfc2460>>.
- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", [RFC 4821](#), DOI 10.17487/RFC4821, March 2007, <<https://www.rfc-editor.org/info/rfc4821>>.
- [RFC6935] Eubanks, M., Chimento, P., and M. Westerlund, "IPv6 and UDP Checksums for Tunneled Packets", [RFC 6935](#), DOI 10.17487/RFC6935, April 2013, <<https://www.rfc-editor.org/info/rfc6935>>.
- [RFC6936] Fairhurst, G. and M. Westerlund, "Applicability Statement for the Use of IPv6 UDP Datagrams with Zero Checksums", [RFC 6936](#), DOI 10.17487/RFC6936, April 2013, <<https://www.rfc-editor.org/info/rfc6936>>.
- [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>>.

10.2. Informative References

- [RFC7510] Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black, "Encapsulating MPLS in UDP", [RFC 7510](#), DOI 10.17487/RFC7510, April 2015, <<https://www.rfc-editor.org/info/rfc7510>>.

Authors' Addresses

Xiaohu Xu
Alibaba Inc.

Email: xiaohu.xxh@alibaba-inc.com

Greg Shepherd
Cisco

Email: gjshep@gmail.com