

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 27, 2017

X. Xu, Ed.
Huawei
S. Somasundaram
Alcatel-Lucent
C. Jacquenet, Ed.
France Telecom
R. Raszuk
Bloomberg LP
Z. Zhang, Ed.
Juniper
September 23, 2016

A Transport-Independent Bit Index Explicit Replication (BIER)
Encapsulation Header
draft-xu-bier-encapsulation-06

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. This document proposes a transport-independent BIER encapsulation header which is applicable regardless of the underlying transport technology.

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 27, 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
1.1.	Requirements Language	2
2.	Terminology	2
3.	BIER Header	3
4.	BIER Header Transport	5
5.	Acknowledgements	5
6.	IANA Considerations	5
7.	Security Considerations	5
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	6
	Authors' Addresses	7

[1.](#) Introduction

Bit Index Explicit Replication (BIER) [[I-D.ietf-bier-architecture](#)] 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 [[I-D.ietf-bier-architecture](#)], BIER adds a header to a multicast data packet (e.g., an IP packet or an MPLS packet). The BIER header carries the information needed for supporting the BIER forwarding procedures. This document proposes a transport-independent BIER encapsulation header which is applicable regardless of the underlying transport technology.

[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 [[I-D.ietf-bier-architecture](#)].

3. BIER Header

The BIER header is shown in Figure 1.

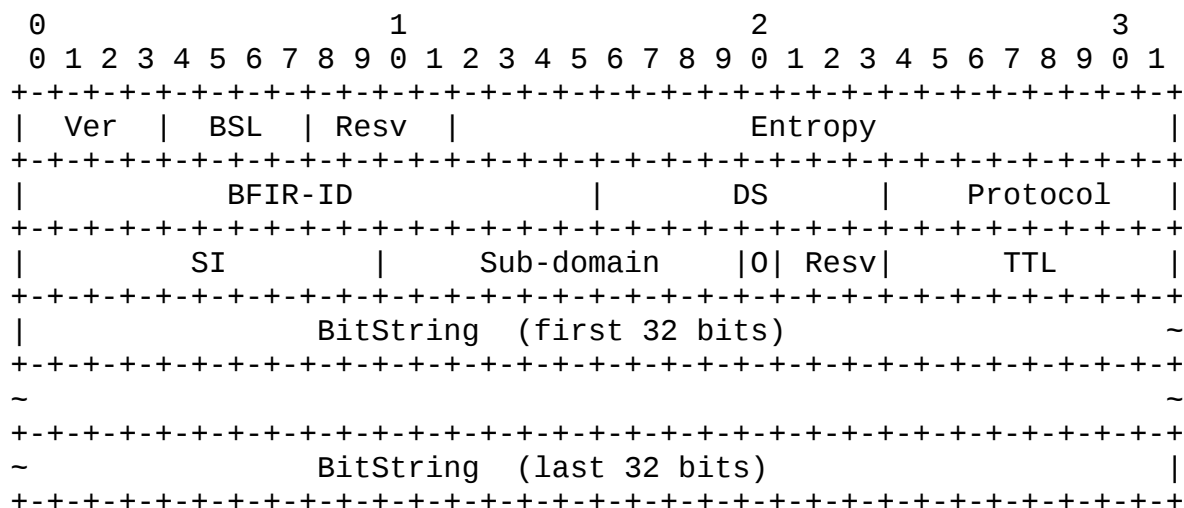


Figure 1: BIER Header Format.

Ver(sion): a 4-bit field identifying the version of the BIER header. This document specifies version 0 of the BIER header.

BSL: Bit String Length. If k is the length of the BitString, the value of this field is $\log_2(k)-5$. However, only the following values are supported [[I-D.ietf-bier-mpls-encapsulation](#)] :

- * 64 bits
- * 128 bits
- * 256 bits
- * 512 bits
- * 1024 bits
- * 2048 bits
- * 4096 bits

The value of the BSL field MUST NOT be set to any value other than those listed above. A received packet containing another value in this field SHOULD be discarded, and an error logged.

Entropy: a 20-bit field containing an "entropy" value that can be used for load balancing purposes.

BFIR-ID: a 2-octet field encoding the BFR-ID of the Bit-Forwarding Ingress Router (BFIR), in the BIER sub-domain where the packet is forwarded to.

DS: The usage of this field is no different from that of the Differentiated Services (DS) field of IPv4 and IPv6 headers. [[RFC2474](#)].

Protocol: a one-octet field indicating the protocol type of the BIER payload as per the protocol numbers used in the Protocol field [[RFC5237](#)] of the IPv4 header and the Next Header field of an IPv6 header. Theoretically, with the use of the IPv4/IPv6 equivalent Protocol/Next-Header field, the payload of a BIER packet could be any type of payload permitted within the IPv4/IPv6 packet. However, this document currently only considers the use of the BIER encapsulation for "transport" type services where the encapsulated packets are either L2.5 (e.g., MPLS) or L3 (e.g., IP) packets. The valid BIER payload types include (but are not limited to) IPv4, IPv6, MPLS, VXLAN [[RFC7348](#)], VXLAN-GPE [[I-D.ietf-nvo3-vxlan-gpe](#)] and LISP [[RFC6830](#)]. The corresponding IP Protocol numbers for VXLAN, VXLAN-GPE and LISP are to be allocated by IANA.

SI: a 10-bit field encoding the Set-Identifier (SI) for this packet in the case where the O-Flag bit is cleared.

Sub-domain: a 10-bit field encoding the sub-domain where the packet is forwarded to in the case where the O-Flag bit is cleared.

O-flag: If it's set, the leftmost 20 bits (i.e., the combination of the SI and Sub-domain fields) should be interpreted as an opaque number signaled and used much as the MPLS-BIER label in MPLS BIER encapsulation [[I-D.ietf-bier-mpls-encapsulation](#)]. A BFR signals the mapping between the opaque number and (subdomain-id, set-id, BSL) and it derives the (subdomain-id, set-id, BSL) from the opaque number in an incoming BIER header just like in the MPLS BIER encapsulation case, except that regular MPLS signaling and forwarding infrastructure is not required.

TTL: The usage of this field is no different from that of the Time To Live (TTL) field in the IPv4 header.

BitString: a variable-length BitString field that, together with the SI and Sub-domain fields, identifies all the destination BFERs for this packet.

4. BIER Header Transport

Since the BIER header format as specified in [Section 3](#) is transport-independent by design, it can be carried with any type of transport encapsulation headers, such as an Ethernet header, a PPP header, an IP header, an MPLS header, a GRE header, an UDP header, etc. Any possible transport encapsulation header must be able to indicate the payload is an BIER header. For instance, in the BIER-in-MAC encapsulation case, the EtherType [[ETYPES](#)] field of the Ethernet header is used for that purpose. In the BIER-in-IP encapsulation case, the Protocol field of the IPv4 header or the Next-Header field of the IPv6 header are used. In the MPLS transport case, the MPLS-BIER encapsulation [[I-D.ietf-bier-mpls-encapsulation](#)] SHOULD be used instead. Note that more details about BIER-in-UDP encapsulation (e.g., UDP checksum and congestion control) will be specified in future versions of this document.

5. Acknowledgements

Thanks Antoni Przygienda, IJsbrand Wijnands, Eric Rosen and Toerless Eckert for their valuable comments and suggestions on this document.

6. IANA Considerations

This document includes a request to IANA to allocate an EtherType code, a PPP protocol code, an IPv4 protocol code/an IPv6 Next-Header code, a UDP destination port to indicate that BIER-encapsulated data follows. Furthermore, this document includes a request to IANA to allocate IP Protocol numbers for VXLAN, VXLAN-GPE and LISP respectively.

7. Security Considerations

As mentioned in [[I-D.ietf-bier-architecture](#)], when BIER is paired with any transport underlay, it inherits the security considerations of the corresponding transport layer. Also, SI and BFIR-ID fields of the BIER header may carry values other than those intended by the BFIR at the risk of misdelivering the packet. Means to protect BFR routers against Man-in-the-Middle and Denial of Service attacks must be provided.

8. References

8.1. Normative References

- [ETYPES] The IEEE Registration Authority, "IEEE 802 Numbers", 2012.
- [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-04](#) (work in progress), July 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

8.2. Informative References

- [I-D.ietf-bier-mpls-encapsulation]
Wijnands, I., Rosen, E., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication in MPLS Networks", [draft-ietf-bier-mpls-encapsulation-05](#) (work in progress), July 2016.
- [I-D.ietf-nvo3-vxlan-gpe]
Kreeger, L. and U. Elzur, "Generic Protocol Extension for VXLAN", [draft-ietf-nvo3-vxlan-gpe-02](#) (work in progress), April 2016.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", [RFC 2474](#), DOI 10.17487/RFC2474, December 1998, <<http://www.rfc-editor.org/info/rfc2474>>.
- [RFC5237] Arkko, J. and S. Bradner, "IANA Allocation Guidelines for the Protocol Field", [BCP 37](#), [RFC 5237](#), DOI 10.17487/RFC5237, February 2008, <<http://www.rfc-editor.org/info/rfc5237>>.
- [RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", [RFC 6830](#), DOI 10.17487/RFC6830, January 2013, <<http://www.rfc-editor.org/info/rfc6830>>.

[RFC7348] Mahalingam, M., Dutt, D., Duda, K., Agarwal, P., Kreeger, L., Sridhar, T., Bursell, M., and C. Wright, "Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", [RFC 7348](https://www.rfc-editor.org/info/rfc7348), DOI 10.17487/RFC7348, August 2014, <<http://www.rfc-editor.org/info/rfc7348>>.

Authors' Addresses

Xiaohu Xu (editor)
Huawei

Email: xuxiaohu@huawei.com

S Somasundaram
Alcatel-Lucent

Email: somasundaram.s@alcatel-lucent.com

Christian Jacquenet (editor)
France Telecom

Email: christian.jacquenet@orange.com

Robert Raszuk
Bloomberg LP

Email: robert@raszuk.net

Zhaohui Zhang (editor)
Juniper

Email: zzhang@juniper.net