Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 28, 2015

X. Xu
Huawei
S. Somasundaram
Alcatel-Lucent
C. Jacquenet
France Telecom
R. Raszuk
Mirantis Inc.

February 24, 2015

BIER Encapsulation
draft-xu-bier-encapsulation-02

Abstract

Bit Index Explicit Replication (BIER) is a new multicast forwarding paradigm which doesn't require an explicit tree-building protocol and doesn't require intermediate routers to maintain any multicast state. This document proposes a transport-independent BIER encapsulation header which is applicable in any kind of transport networks.

Table of Contents

1. Introduction

Bit Index Explicit Replication (BIER) [I-D.wijnands-bier-architecture] is a new multicast forwarding paradigm which doesn't require an explicit tree-building protocol and doesn't require intermediate routers to maintain any multicast state. As described in [I-D.wijnands-bier-architecture], BIER requires that a multicast data packet (e.g., an IP packet or an MPLS packet) to be encapsulated with a BIER header that carries the information needed for supporting the BIER forwarding procedures. This information at least includes Set-Identifier (SI), Multi-Topology Identifier (MT-ID) and BitString. The SI and the BitString are used together to identify the set of egress BFRs (BFERs) to which the packet must be delivered. In addition, to indicate what type of payload is following the BIER header, a protocol type field is neccessary. This document proposes a transport-independent BIER encapsulation header which is applicable in any kind of transport networks.

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.wijnands-bier-architecture].

3. BIER Header

The BIER header is shown as follows:

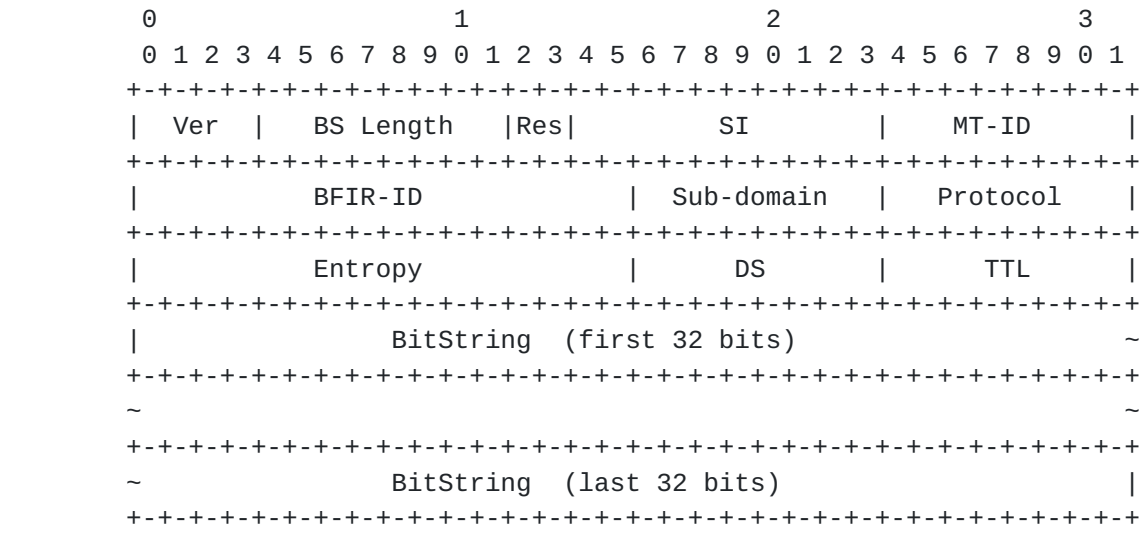

Ver(sion): a 4-bit field identifying the version of the BIER header. This document specifies version 0 of the BIER header.

BS Length: a one-octet field indicating the length of the BitString in 4-byte. Note that legal BS Length values are specified in [I-D.wijnands-bier-architecture].

Res: a 2-bit reserved field.

SI: a 10-bit field encoding the Set-Identifier (SI) for this packet.

MT-ID: a one-octet field indicating which routing topology [RFC4915] [RFC5120] should be applied for BIER forwarding.

BFIR-ID: a 2-octet field encoding the BFR-ID of the BFIR, in the sub-domain to which the packet has been assigned.

Sub-domain: a one-octet field encoding the sub-domain to which the packet has been assigned.

Protocol: a one-octet field indicating the protocol type of the BIER payload as per IP protocol numbers used in the Protocol field

of the IPv4 header and the Next Header field of IPv6 header. The valid BIER payload types include but not limited to IPv4, IPv6, MPLS, VXLAN [RFC7348],VXLAN-GPE [I-D.quinn-vxlan-gpe] , and etc. The corresponding IP Protocol numbers for VXLAN and VXLAN-GPE are to be allocated by IANA.

Entropy: a 2-octet field containing an "entropy" value that can be used for load balancing purposes.

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

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

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

4. Transport Encapsulation for BIER Header

Since the BIER encapsulation format as specified in Section 3 is transport-independent, it can be encapsulated with any type of transport encapsulation headers, such as Ethernet header, PPP header, IP header, MPLS header, GRE header, UDP header etc. It requires for each possible transport encapsulation header to be able to indicate the payload is an BIER header. For instance, In the BIER-in-MAC encapsulation case, the EtherType field in the Ethernet header is used. In the BIER-in-IP encapsulation case, the Protocol field in the IPv4 or or the Next-Header field in the IPv6 header is used. In the BIER-in-MPLS encapsulation case, either the Protocol Type field [I-D.xu-mpls-payload-protocol-identifier] within the MPLS packet or a to-be-assigned Extended Spcial Purpose label [RFC7274] is used.

5. Acknowledgements

TBD.

6. IANA Considerations

This document includes a request to IANA to allocate an EtherType code,a PPP protocol code, an IPv4 protocol code (i.e., an IPv6 Next-Header code), a UDP destination port for carring the BIER-encapsulated packet over the corresponding transport networks. Furthermore, This document includes a request to IANA to allocate IP Protocol numbers for VXLAN and VXLAN-GPE respectively.

7. Security Considerations

TBD.

8. References

8.1. Normative References

[ETYPES] The IEEE Registration Authority, "IEEE 802 Numbers", 2012.

[I-D.wijnands-bier-architecture]
Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", draft-wijnands-bier-architecture-04 (work in progress), February 2015.

[I-D.xu-mpls-payload-protocol-identifier]
Xu, X. and M. Chen, "MPLS Payload Protocol Identifier", draft-xu-mpls-payload-protocol-identifier-00 (work in progress), September 2013.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC7274] Kompella, K., Andersson, L., and A. Farrel, "Allocating and Retiring Special-Purpose MPLS Labels", RFC 7274, June 2014.

8.2. Informative References

[I-D.quinn-vxlan-gpe]
Quinn, P., Agarwal, P., Fernando, R., Lewis, D., Kreeger, L., Smith, M., Yadav, N., Yong, L., Xu, X., Elzur, U., and P. Garg, "Generic Protocol Extension for VXLAN", draft-quinn-vxlan-gpe-03 (work in progress), July 2014.

[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, December 1998.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, June 2007.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.

[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, August 2014.

Authors' Addresses

Xiaohu Xu
Huawei

Email: xuxiaohu@huawei.com

S Somasundaram
Alcatel-Lucent

Email: somasundaram.s@alcatel-lucent.com

Christian Jacquenet
France Telecom

Email: christian.jacquenet@orange.com

Robert Raszuk
Mirantis Inc.

Email: robert@raszuk.net