

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

X. Xu
M. Chen
Huawei
K. Patel
I. Wijnands
Cisco
A. Przygienda
Ericsson
May 27, 2015

BGP Extensions for BIER
draft-xu-idr-bier-extensions-02

Abstract

Bit Index Explicit Replication (BIER) is a new multicast forwarding architecture which doesn't require an explicit tree-building protocol and doesn't require intermediate routers to maintain any multicast state. BIER is applicable in a multi-tenant data center network environment for efficient delivery of Broadcast, Unknown-unicast and Multicast (BUM) traffic while eliminating the need for maintaining a huge amount of multicast state in the underlay. This document describes BGP extensions for advertising the BIER-specific information. These extensions are applicable in those multi-tenant data centers where BGP instead of IGP is deployed as an underlay for network reachability advertisement. These extensions may also be applicable in other scenarios.

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 November 28, 2015.

Internet-Draft

BGP Extensions for BIER

May 2015

Copyright Notice

Copyright (c) 2015 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	3
2.	Terminology	3
3.	BIER Path Attribute	3
4.	Originating BIER Attribute	4
5.	Restrictions on Sending/Receiving	5
6.	Deployment Considerations	5
7.	Acknowledgements	5
8.	IANA Considerations	6
9.	Security Considerations	6
10.	References	6
10.1.	Normative References	6
10.2.	Informative References	6
	Authors' Addresses	7

[1.](#) Introduction


Bit Index Explicit Replication (BIER)

[\[I-D.wijnands-bier-architecture\]](#) is a new multicast forwarding architecture which doesn't require an explicit tree-building protocol and doesn't require intermediate routers to maintain any multicast state. BIER is applicable in a multi-tenant data center network environment for efficient delivery of Broadcast, Unknown-unicast and Multicast (BUM) traffic while eliminating the need for maintaining a huge amount of multicast state in the underlay

[\[I-D.kumar-bier-use-cases\]](#). This document describes BGP extensions

[Page 2]

May 2015






Figure 1:BIER TLV

Type: Two octets encoding the BIER TLV Type: TBD.

Length: Two octets encoding the length in octets of the TLV, including the type and length fields. The length is encoded as an

unsigned binary integer. (Note that the minimum length is 8, indicating that no sub-TLV is present.)

Sub-domain: a one-octet field encoding the sub-domain ID corresponding to the BFR-ID.

BFR-ID: a two-octet field encoding the BFR-ID.

BSL: a one-octet field indicating the length of the Bitstring in 4-octets. The field MUST be filled with one of the valid BSL values as specified in [[I-D.wijnands-bier-architecture](#)]. Upon receiving a BSL-TLV containing an invalid BSL value, it MUST be ignored.

Sub-TLVs: contains one or more sub-TLV. The BIER MPLS Encapsulation sub-TLV is one of such sub-TLVs.

The BIER MPLS Encapsulation sub-TLV is encoded as follows:




Figure 2:BIER MPLS Encapsulation sub-TLV

Type:TBD

Length:8

Label Range Size: a one-octet field indicating the size of the label range.

Label Range Base: a 3-octet field where the 20 rightmost bits represent the first label in the label range while the other bits MUST be set to 0 when transmitting, and MUST be ignored upon receipt.

[4.](#) Originating BIER Attribute

An implementation that supports the BIER attribute MUST support a policy to enable or disable the creation of the BIER attribute and its attachment to specific BGP routes. An implementation MAY disable the creation of the BIER attribute unless explicitly configured to do so otherwise. A BGP speaker MUST only attach the locally created

Xu, et al.

Expires November 28, 2015

[Page 4]

Internet-Draft

BGP Extensions for BIER

May 2015

BIER attribute to a BGP UPDATE message in which at least one of its BFR-prefixes is contained in the NLRI.

[5.](#) Restrictions on Sending/Receiving

An implementation that supports the BIER attribute MUST support a per-EBGP-session policy, that indicates whether the attribute is enabled or disabled for use on that session. The BIER attribute MUST NOT be sent on any EBGP peers for which the session policy is not configured. If an BIER attribute is received on a BGP session for which session policy is not configured, then the received attribute MUST be treated exactly as if it were an unrecognised non-transitive attribute. That is, "it MUST be quietly ignored and not passed along to other BGP peers".

To prevent the BIER attribute from "leaking out" of an BIER domain, each BGP router on the BIER domain MUST support an outbound route announcement policy. Such a policy MUST be disabled on each EBGP session by default unless explicitly configured.

[6.](#) Deployment Considerations

It's assumed by this document that the BIER domain is aligned with the Administrative Domain (AD) which are composed of multiple ASes (either private or public ASes). Use of the BIER attribute in other scenarios is outside the scope of this document.

Since the BIER attribute is an optional, transitive BGP path attribute, a non-BFR BGP speakers could still advertise the received route with a BIER attribute. This is desirable in the incremental deployment scenario where a BGP speaker could tunnel a BIER packet or the payload of a BIER packet to a BFER directly if the BGP next-hop of the route for that BFER is a non-BFR. Furthermore, a BGP speaker is allowed to tunnel a BIER packet to the BGP next-hop if these two BFR-capable BGP neighbors are not directly connected (e.g., multi-hop EBGP) . As for which tunnel type should be used, it could be manually configured or dynamically negotiated by using the BGP Encapsulation SAFI mechanism as defined in [[RFC5512](#)]. The BIER-specific extensions to the BGP Encapsulation SAFI would be defined in a future version of this document.

[7.](#) Acknowledgements

Thanks a lot for Eric Rosen for his valuable comments on this document.

[8.](#) IANA Considerations

IANA is requested to assign a codepoint in the "BGP Path Attributes" registry to the BIER attribute. IANA shall create a registry for "BGP BIER Attribute Types". The type field consists of two octets, with possible values from 1 to 65535 (The value 0 is "reserved".) The allocation policy for this field is to be "First Come First Serve". Type codes should be allocated for BIER TLV and BIER MPLS Encapsulation sub-TLV respectively.

[9.](#) Security Considerations

This document introduces no new security considerations beyond those already specified in [[RFC4271](#)].

[10.](#) References

[10.1.](#) Normative References

- [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-05](#) (work in progress), March 2015.
- [I-D.wijnands-mpls-bier-encapsulation]
Wijnands, I., Rosen, E., Dolganow, A., Tantsura, J., and S. Aldrin, "Encapsulation for Bit Index Explicit Replication in MPLS Networks", [draft-wijnands-mpls-bier-encapsulation-02](#) (work in progress), December 2014.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC5512] Mohapatra, P. and E. Rosen, "The BGP Encapsulation Subsequent Address Family Identifier (SAFI) and the BGP Tunnel Encapsulation Attribute", [RFC 5512](#), April 2009.

[10.2.](#) Informative References

- [I-D.ietf-rtgwg-bgp-routing-large-dc]
Lapukhov, P., Premji, A., and J. Mitchell, "Use of BGP for routing in large-scale data centers", [draft-ietf-rtgwg-bgp-routing-large-dc-02](#) (work in progress), April 2015.

Xu, et al. Expires November 28, 2015 [Page 6]

Internet-Draft BGP Extensions for BIER May 2015

- [I-D.kumar-bier-use-cases]
Kumar, N., Asati, R., Chen, M., Xu, X., Dolganow, A., Przygienda, T., arkadiy.gulko@thomsonreuters.com, a., and D. Robinson, "BIER Use Cases", [draft-kumar-bier-use-cases-02](#) (work in progress), February 2015.

Authors' Addresses

Xiaohu Xu
Huawei

Email: xuxiaohu@huawei.com

Mach Chen
Huawei

Email: mach.chen@huawei.com

Keyur Patel
Cisco

Email: keyupate@cisco.com

IJsbrand Wijnands
Cisco

Email: ice@cisco.com

Antoni Przygienda
Ericsson

Email: antoni.przygienda@ericsson.com