Workgroup: BIER
Internet-Draft:
draft-ietf-bier-multicast-as-a-service-03
Published: 6 February 2024
Intended Status: Standards Track
Expires: 9 August 2024
Authors: Z. Zhang (Juniper Networks)
E. Rosen (Individual)
D. Awduche (Individual)
G. Shepherd (Individual)
Z. Zhang (ZTE Corporation)
G. Mishra (Verizon)

Multicast/BIER As A Service

Abstract

This document describes a framework for providing multicast as a service via Bit Index Explicit Replication (BIER) [RFC7279], and specifies a few enhancements to [draft-ietf-bier-idr-extensions] [RFC8279] [RFC8401] [RFC8444] to enable multicast/BIER as a service.

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 RFC2119.

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 9 August 2024.

Copyright Notice

Table of Contents

1. Introduction

Currently multicast is primarily used in the following scenarios:

- *Enterprise Applications. For example, large scale financial data publishing.

- *Provider/underlay tunnels for MVPN and for EVPN BUM.

- *Real-time IPTV offered by a service provider to its customers.

Besides the above, large scale multicast services, especially transit multicast transport provided by large Internet Service Providers is virtually non-existent. This is mainly because of the following chicken and egg dilemma:

- *Traditional multicast technologies are complicated and lack scalability. The revenue that multicast services bring in cannot

offset the Capex and Opex that an operator has to invest, so provider networks typically do not enable multicast even though the deployed equipment does support multicast.

*As a result, Content Providers cannot take advantage of multicast and instead use less efficient methods like Ingress Replication, Peer2Peer, or multicast at application layer.

A recent multicast technology breakthrough, BIER, provides a simple and scalable solution for large scale multicast deployment, independent of number of multicast flows. In the meantime, large scale distribution of ultra high definition video content has become more and more popular and important. Service providers simply cannot keep on increasing their network capacity even if they could shift cost to Content Providers. With these developments, service providers now have both the need and means to provide scalable multicast service, potentially across multiple providers.

This document describes a framework for Multicast As A Service (MAAS) enabled by BIER. We use Content Delivery Network (CDN) as example, though it applies to any large scale multicast delivery service.

1.1. Terminologies

Readers are assumed to be familiar with multicast, BIER, BGP and ISIS/OSPF concepts and procedures. Some terminologies are listed here for convenience.

*BFR: BIER Forwarding Router.

*BFIR: BIER Forwarding Ingress Router.

*BFER: BIER Forwarding Egress Router.

*EBFR: Edge BFR. Including BFIR and BFER.

*BSL: BitStrengLength. Number of bits in the BitString of a BIER header.

1.2. A CDN of A Single Provider

To make it easier to understand, we first consider a simple example: a CDN owned by a single operator, which could be a Content Provider itself. The network spans multiple ASes as shown in the following figure:

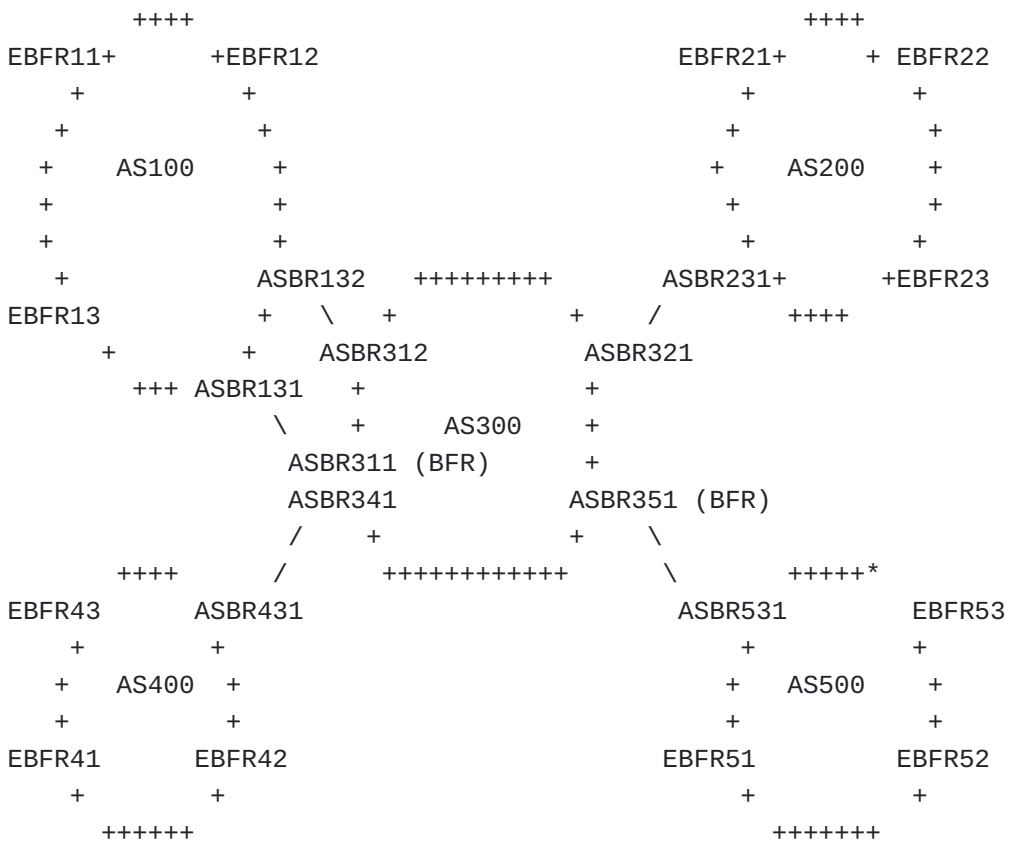

The CDN uses BIER for multicast transport and Edge BIER Forwarding Routers (EBFRs) are located throughout the network. Some of them are connected towards multicast content sources and are referred to as BIER Forwarding Ingress Routers (BFIRs) in BIER architecture. Most of them are connected towards multicast content receivers and are referred to as BIER Forwarding Egress Routers (BFERs). Notice that between content sources and BFIRs there may be Protocol Independent Multicast (PIM) in use, while between content receivers and BFERs there may be PIM and/or IGMP in use.

At the initial deployment stage, there might be only a few transit BIER Forwarding Routers (BFRs) at strategic points in the network (e.g. ASBR311 and ASBR351). BGP sessions are established among the EBFRs and BFRs, and BGP extensions as defined in [I-D.ietf-bier-idr-extensions] are used to signal BIER information. All these are in a single BIER sub-domain.

In the example of initial stage with only ASBR311 and ASBR351 as BFRs, multicast traffic arriving at EBFR11 will be imposed with a BIER header and replicated to EBFR12/EBFR13/ASBR311 over tunnels. ASBR311 will further replicate traffic to ASBR351/EBFR41/EBFR42/ EBFR43/EBFR21/EBFR22/EBFR23 over tunnels, and ASBR351 will further replicate traffic to EBFR51/EBFR52/EBFR53 over tunnels.

The BGP signaling can be explained using the following example. EBFR43 advertises its BFR-prefix (a loopback address) as /32 IPv4 or /128 IPv6 prefix in BGP with a BIER Path Attribute (BPA) [I-D.ietf-bier-idr-extensions]. The BIER Nexthop sub-TLV is set to the same BFR-prefix. ASBR431 receives it and re-advertises it but does not do anything wrt BIER or the BPA because it does not support BIER. Same happens on ASBR341. When ASBR311 and ASBR351 receive it from ASBR341, they create a BIFT entry corresponding to EBFR43's BFR-ID. The entry causes a BIER packet with corresponding bit set in its BitString to be tunneled to EBFR43 (the BIER Nexthop in the BPA). When ASBR311 and ASBR351 re-advertise EBFR43's BFR-prefix, they update the BIER Nexthop in the BPA to their own BFR-prefix, so that when eventually EBFR11 receives the re-advertised route, it creates a BIFT entry that causes corresponding packets to be tunneled to ASBR311 (the BIER Nexthop in the BPA).

Over time, more routers in network may be upgraded to support BIER and become a BFR. For example, once ASBR431 is upgraded to a BFR, ASBR311 no longer needs to tunnel traffic to EBFR41/EBFR42/EBFR43 but only need to tunnel one copy to ASBR431, who will then replicate to EBFR41/EBFR42/EBFR43.

1.2.1. IGP/BGP Interworking

Additionally, if enough routers in an AS (or just one of its IGP areas) can be upgraded to run BIER, then hop-by-hop BIER forwarding can be utilized there, using IGP extensions for BIER signaling [RFC8401] [RFC8444].

Notice that even with this there is still only one BIER sub-domain, with mixed IGP and BGP signaling for BIER. To redistribute BIER information between IGP and BGP, procedures specified in [I-D.ietf-bier-prefix-redistribute] and detailed in Section 2.2 are followed.

1.3. A CDN That Involves Another Provider

In the above example, the CDN is providing multicast transport service, with simplicity and scalability provided by BIER (the per-flow state is confined to the edges). Now let us go one step further and consider that AS300 belongs to a different Internet Service Provider. Now the ISP is providing BIER As A Service (BAAS) to the CDN, by being part of the CDN's BIER sub-domain. Notice that, not only does the ISP not have per-tree state (it does not have EBFRs), but also its BFRs do not need BFR-ID assigned. The ISP does need to learn about all the EBFRs and their corresponding BFR-IDs (through signaling).

1.3.1. Providing Independent BAAS To Multiple Customers

Now consider that the ISP also provides BAAS for another CDN. Each of the two CDNs has its own BIER domain, with its own BFR-ID or even sub-domain ID assignment that could conflict between the two CDNs. For example, both have BFR-ID 100 and sub-domain ID 0 assigned but they are totally independent of each other. For an BFR in the ISP to support this, with BGP signaling it needs to advertise its own BFR-prefix multiple times, each time with a different RD that is mapped to the corresponding CDN. A new SAFI BIER (to be allocated by IANA) is used.

In the above example, there are two paths between AS100 and AS300. It is possible that while ASBR311 is the BFR, ASBR312 is the unicast best path into AS300 and beyond from AS100. Advertising BFR-prefixes using a different SAFI with a RD also has the side benefit of allowing incongruent topologies for unicast and BIER.

In the existing BIER architecture and IGP extensions for BIER a sub-domain is tied to a single topology (either the one and only topology if Multi-topology ISIS/OSPF is not used, or a topology as defined in Multi-topology ISIS/OSPF). In the BIER sub-TLV that ISIS/OSPF attaches to a BFR-prefix, a Sub-domain-ID value can only appear once for a particular topology. In this document, a BFR in the BAAS provider may belong to different and independent BIER domains, and the same sub-domain ID needs to be signaled multiple times, once for each BIER domain (notice that the same sub-domain-ID actually identifies different sub-domains in different BIER domains, so this does not really change the architectural requirement that a sub-domain is tied to a single topology). To do so, a new "BIER Domain" sub-TLV is introduced, and its value field includes a RD (as in the BGP signaling) and a BIER sub-sub-TLV that is the same as currently specified in ISIS/OSPF extensions for BIER.

This works very well because of the flexible BIER architecture - a BIER packet is forwarded based on a Bit Index Forwarding Table (BIFT) that is determined by a 20-bit BIFT ID in front of the BIER header, and each (subdomain, BSL, set) tuple has its own BIFT. Traditionally, a subdomain is identified by a sub-domain ID but in this document a subdomain is now identified by a (RD, sub-domain ID) tuple in the control plane.

With this, the scaling aspect on a BFR comes to how many BAAS customer the provider needs to support. For example, if it needs to support 16 BAAS customers, one BSL, and four sets (Section 1.4.1) for each customer, then the provider needs to support 64 BIFTs (16 x 1 x 4). If the BSL is 256, then each BIFT has 256 entries in it and the total number of BIFT entries (routes) is 4k (256 x 64). Notice that this 4k number is not related to the number of customers'

multicast flows, but only related to the number of customers and number of customer EBFRs. The number of customers with their own independent BIER domains are likely not very large initially, but if multicast as a service gets more widely used, the protocol and procedures defined in this document can scale up to the extent of how many BIFTs (and BIFT entries) a BFR can support. Since there is no real difference between a BIFT entry and a unicast RIB/FIB entry, as long as the scaling requirements are adequately considered in the BIER forwarding plane implementation (e.g., enough memory is allocated for the BIFTs), scaling will not become a bottleneck.

Building/updating the BIFTs is the same as in the base BIER architecture, except that in the control plane a subdomain is identified by a (RD, sub-domain ID) tuple instead of just a sub-domain ID. This is transparent to the forwarding plane - a BIFT is always identified by an opaque 20-bit opaque number. This opaque number is either a label for MPLS encapsulation or an opaque number for non-MPLS encapsulation, and the optional static encoding as specified in [I-D.ietf-bier-non-mpls-bift-encoding] cannot be used.

1.3.2. BIER VPN

In the previous section, customer BFR-prefixes (with non-zero BFR-IDs) are advertised into the provider network with BIER information. Notice that it is not a VPN scenario - even though the customers and the provider may have independent and overlapped BIER sub-domain and BFR-IDs, the customer networks and the provider network are in the same address space.

While the same mechanism can be used for VPN, this section looks at VPN scenario at a different angle and considers more aspects. It will become clear that these apply to non-VPN scenario as well.

1.3.2.1. Signaling in/over Provider Underlay

Advertising non-VPN customer BFR-prefixes into provider network as described in Section 1.3.1 may be considered acceptible (this is just like advertising internet routes), but advertising VPN customer BFR-prefixes into provider network is not aligned with VPN architecture [RFC4364].

Notice that while BIER signaling attaches BIER information to BFR-prefixes that are used to find routing underlay paths to BFERs, it is the BFER-IDs that are used in BIFTs. In other words, the key is to be able to find paths to BFERs that are identified as BFR-IDs (or bit positions in BIER packets' BitString to be more precise). Since customers BFERs will be reached via PEs, as long as P routers know which customer BFR-IDs are associted with (i.e., reachable via) which PEs, they can calculate the BIFTs for the customer BIER

domains without knowing customer BFR-prefixes. This association information can be advertised by PEs via Proxy range sub-TLV [I-D.ietf-bier-prefix-redistribute].

The signaling discussed above is in the provider underlay, with customer BFR-IDs advertised with PE BFR-prefixes. In pararell, with overlay signaling via MPLS-labeled VPN address SAFI [RFC4364] among PEs, BIER Path Attribute [I-D.ietf-bier-idr-extensions] carries customer BIER information (including customer BFR-prefixes) from site to site, with redistribution [I-D.ietf-bier-prefix-redistribute] between IGP and BGP at PEs.

The underlay/overlay signaling is illustrated in the following figure. Customer sites run customer IGP/BGP towards the VRFs on PEs, with customer BIER information advertised with customer BFR-prefixes. Overlay signaling then happens among PEs, with the complete BIER information carried in BIER Path Attribute attached to cutomer BFR-prefixes (that are advertised as normal VPN routes), and redistribution happens at the PEs. In the mean time, underlay signaling among PE/P routers advertises customer BFR-ID and sub-domain information via BIER proxy range sub-TLVs.

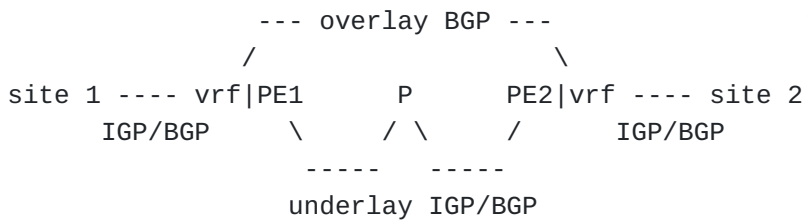

1.3.2.2. VPN vs. Non-VPN

As alluded to earlier, the VPN procedures (in particular using BIER Proxy Range sub-TLV) can be used for non-VPN scenarios as well, and the overlay signaling for the non-VPN case can be via BGP-LU [RFC3107] with BIER Path Attibute attached.

1.3.3. Control and Accounting

With BGP based signaling, internal routers of a BAAS provider does not need explicit configuration for the BIER transport services that it support. In the above example, the ASBRs (ASBR311, ASBR312, ASBR321, ASBR341, ASBR351) in AS300 only need to have BGP policy configured to allow certain received BFR-prefix advertisements to trigger necessary BIER state and additional signaling of their own. For example, when ASBR351 receives the BFR-prefix advertisement, if its local configuration allows it may create corresponding BIFTs and BIFT entries, and additionally originates or updates its own BFR-prefix advertisement. An internal BFR inside AS300, upon receiving

the BGP advertisements, may or may not need to go through the same policy check again (based on the providers operation model).

When the ASBRs (re-)advertise BFR-prefixes toward their external peers, they could enable statistics counters for the corresponding BIER labels so that they can count incoming BIER packets from external peers specifically for this BAAS. Similarly, the ASBRs can enable statistics counters for BIER labels they receive from external peers, so that they can count outgoing BIER packets delivered to the external peers. These incoming and outgoing counters can be used for accounting and billing purposes.

1.4. Sets and Segmentation

The number of EBFRs could very well be larger than the BSL. There are two ways to handle that - multiple sets or segmentation.

1.4.1. Multiple Sets

With this method the set of EBFRs are grouped into multiple sets, and the number of EBFRs in a set is smaller than the BSL. A BFIR may need to send multiple copies of a multicast packet to reach all BFERs, one copy for each set that covers one or more expecting BFERs. A separate BIFT is needed for each set (because the same bit in the BitString of packets for different sets maps to different BFERs). This not only leads to multiple copies to be sent over the same link, but also requires additional BIFTs. In the earlier example, 64 BIFTs are needed for 16 BAAS customers because each customer needs 4 BIFTs for the multiple sets.

1.4.2. Segmentation

With this method, a BIER network is segmented into multiple regions, each with its own BIER sub-domain. In the earlier example, each AS could be an independent sub-domain. A BIER packet from EBFR11 will be decapsulated by the segmentation border router ASBR311, and then sent into next sub-domain in AS300 with a new BIER header. The segmentation [RFC7524] involves Multicast Flow Overlay [RFC8279] [RFC8556] so that the segmentation border routers know what BitString to use when sending onto the next segment. The advantage of segmentation is that only a single copy needs to be sent, and the number of BIFTs is also reduced on all BFRs. The disadvantage is that the segmentation points need to run multicast flow overlay protocol and maintain related state in control plane and data plane.

A deployment may start without the need for either multiple sets or segmentation when the number of EBFRs is small. When the number of EBFRs grows, segmentation can be introduced incrementally. A new BFR can be added as, or an existing BFR could be converted to, a segmentation point, splitting the original sub-domain into two

independent sub-domains. The segmentation point does not re-advertise BIER information from one sub-domain to another. Other BFRs/EBFRs do not need any configuration changes except to make sure that all BIER information exchange is restricted to a single sub-domain (for example, two BFRs were BGP peers before and were exchanging BIER information but now they belong to two sub-domains and only exchange BIER information with the segmentation point and other BFRs in the same sub-domain).

In the earlier example of a CDN of a single provider, using segmentation may be acceptable, even though the overlay state needs to be kept by the segmentation points. A BAAS provider may need to carefully consider if it wants to keep a customer's overlay state on those segmentation points. On the other hand, the provider may consider hosting per-customer segmentation points. For example, tethering small or virtual BFRs to an ASBR and have those BFRs be the segmentation points [I-D.ietf-bier-tether].

2. Specifications for Enhancements to BIER Signaling with BGP/IGP

2.1. BGP Procedures

A BFR receiving the advertisement MUST use the tunnel destination in the TEA to determine where to forward a BIER packet whose BitString has a set bit corresponding to the BFR-prefix, unless the TEA does not exist, in which case the BFR-prefix itself is used for the determination. When the BFR re-advertises the BFR-prefixes, it MUST change the tunnel destination in the TEA to itself, or add a TEA with the tunnel destination set to itself if there was no TEA in the received advertisement.

The TEA SHOULD have a Protocol Sub-TLV with protocol type BIER (0xAB37).

A transit BFR that is allowed (by provisioning or based on policy) to participate in a BIER sub-domain MUST advertise its own BFR-prefix with a BPA. The BFR-id in the BPA SHOULD be 0. Depending on the operational model of the operator, the advertisement MAY be based on received BFR-prefixes (subject to certain BGP policy verification), or MAY do so only with explicit configuration.

If a provider provides independent BAAS services to multiple customers, when its BFR receives BFR-prefixes from a customer it MUST re-advetise with a new BIER SAFI. For simplicity, all BFRs of the provider use the same RD that is specifically assigned for the customer. When a BFR re-advertises BFR-prefixes to a customer, it MUST re-advertise with SAFI 1 or 2.

If multiple providers together provide BAAS to a customer, then the two providers may assign the same RD for the customer or do RD

rewriting when re-advertising BFR-prefixes from one provider to another.

2.2. ISIS/OSPF Procedures

This document defines a new BIER Domain Sub-TLV of ISIS TLVs 135, 235, 236, and 237. The sub-TLV type is to be allocated.

This document also defines a new BIER Domain Sub-TLV of OSPF Extended Prefix TLV. The sub-TLV type is to be allocated.

The value part of the BIER Domain Sub-TLV includes a 64-bit Route Distinguisher followed by one or more BIER Info Sub-TLV (as defined in [RFC8401] and [RFC8444] respectively) as its sub-sub-TLVs .

When a BFR redistribute a BFR-prefix from BGP into ISIS/OSPF, if the BGP advertisement is of BIER SAFI, a BIER Domain sub-TLV is attached, with the RD part of the sub-TLV copied from the BGP advertisement. For each BIER TLV in the BPA, a BIER Info sub-sub-TLV is added in the BIER Domain sub-TLV, with the subdomain-id and BFR-id copied from the corresponding BIER TLV in the BPA, and the Encapsulation sub-sub-sub-TLV omitted because it is not needed.

If the BGP advertisement is of SAFI 1 or 2, BIER Info Sub-TLVs are constructed as above directly, without using a BIER Domain sub-TLV.

When a BFR redistribute a BFR-prefix from ISIS/OSPF into BGP, if there is a BIER Domain sub-TLV in the corresponding ISIS LSP or OSPF LSA, the BGP advertisement is of BIER SAFI and the RD part of the NLRI is set to the RD from the BIER Domain sub-TLV. For each BIER Info sub-sub-TLV in the BIER Domain sub-TLV, a BIER TLV is included in the BPA, with the subdomain-id and BFR-id copied from the corresponding BIER Info sub-sub-TLV. The MPLS Encapsulation sub-TLV is omitted. The tunnel destination in the TEA is set to the BFR's BFR-prefix.

If there is no BIER Domain sub-TLV in the corresponding ISIS LSP or OSPF LSA for the BFR-prefix, the BGP advertisement is of SAFI 1 or 2, and the BPA is constructed similar to the above (the only difference is that in this case BIER Info sub-TLVs are not part of a BIER Domain sub-TLV).

3. IANA Considerations

This document requests the following IANA assignments:

* A sub-TLV type for BIER Domain Sub-TLV from ISIS "Sub-TLVs for TLVs 135, 235, 236, and 237" registry.

*A sub-TLV type for BIER Domain Sub-TLV from OSPFv2 Extended Prefix Sub-TLV registry.

*A BIER SAFI from Subsequent Address Family Identifiers (SAFI) registry.

4. Security Considerations

There are no security concerns wrt exchange of BIER information besides what have been discussed in [I-D.ietf-bier-idr-extensions] and [RFC8401] [RFC8444].

The tunnels between BFRs that are not directly connected are ideally auto-configured to reduce provisioning burdens. Given that they may span multiple ASes and MPLS may not always be available, BIER over UDP/GRE/IPv4/IPv6 becomes very convenient, though that has the same security concerns well discussed in "Security Considerations" of [RFC4023] and [RFC7510].

As one mitigation when the tunnel is not secured, a BFR MAY use source address filtering based on pre-provisioned or dynamically learned allowable addresses. With dynamic learning, if a BFR receives a BFR-prefix with a BPA and a TEA (see Section 2.1), it sets up a forwarding filter to allow IP/GRE/UDP tunneling from the address encoded in the "Tunnel Egress Endpoint" sub-TLV of Tunnel TLVs in the TEA. While that is the address for this BFR to tunnel traffic to, this BFR will also likely receive tunneled traffic from that address.

5. Acknowledgements

The authors thank Lenny Giuliano and Antoni Przygenda for their review and suggestions.

6. References

6.1. Normative References

[I-D.ietf-bier-idr-extensions]
Xu, X., Chen, M., Patel, K., Wijnands, I., Przygienda, T., and Z. J. Zhang, "BGP Extensions for BIER", Work in Progress, Internet-Draft, draft-ietf-bier-idr-extensions-10, 13 June 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-bier-idr-extensions-10>.

[I-D.ietf-bier-prefix-redistribute]
Zhang, Z., Wu, B., Zhang, Z. J., Wijnands, I., Liu, Y., and H. Bidgoli, "BIER Prefix Redistribute", Work in Progress, Internet-Draft, draft-ietf-bier-prefix-

redistribute-05, 7 September 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-bier-prefix-redistribute-05>.

[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>.

[RFC8401] Ginsberg, L., Ed., Przygienda, T., Aldrin, S., and Z. Zhang, "Bit Index Explicit Replication (BIER) Support via IS-IS", RFC 8401, DOI 10.17487/RFC8401, June 2018, <https://www.rfc-editor.org/info/rfc8401>.

[RFC8444] Psenak, P., Ed., Kumar, N., Wijnands, IJ., Dolganow, A., Przygienda, T., Zhang, J., and S. Aldrin, "OSPFv2 Extensions for Bit Index Explicit Replication (BIER)", RFC 8444, DOI 10.17487/RFC8444, November 2018, <https://www.rfc-editor.org/info/rfc8444>.

[RFC8556] Rosen, E., Ed., Sivakumar, M., Przygienda, T., Aldrin, S., and A. Dolganow, "Multicast VPN Using Bit Index Explicit Replication (BIER)", RFC 8556, DOI 10.17487/RFC8556, April 2019, <https://www.rfc-editor.org/info/rfc8556>.

6.2. Informative References

[I-D.ietf-bier-non-mpls-bift-encoding] Wijnands, I., Mishra, M. P., Xu, X., and H. Bidgoli, "An Optional Encoding of the BIFT-id Field in the non-MPLS BIER Encapsulation", Work in Progress, Internet-Draft, draft-ietf-bier-non-mpls-bift-encoding-04, 30 May 2021, <https://datatracker.ietf.org/doc/html/draft-ietf-bier-non-mpls-bift-encoding-04>.

[I-D.ietf-bier-tether] Zhang, Z. J., Warnke, N., Wijnands, I., and D. O. Awduche, "Tethering A BIER Router To A BIER incapable Router", Work in Progress, Internet-Draft, draft-ietf-bier-tether-04, 5 July 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-bier-tether-04>.

[RFC7524] Rekhter, Y., Rosen, E., Aggarwal, R., Morin, T., Grosclaude, I., Leymann, N., and S. Saad, "Inter-Area Point-to-Multipoint (P2MP) Segmented Label Switched Paths (LSPs)", RFC 7524, DOI 10.17487/RFC7524, May 2015, <https://www.rfc-editor.org/info/rfc7524>.

[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>.

Authors' Addresses

Zhaohui Zhang
Juniper Networks

Email: zzhang@juniper.net

Eric Rosen
Individual

Email: erosen52@gmail.com

Daniel Awduche
Individual

Email: awduche@awduche.com

Greg Shepherd
Individual

Email: gjshep@gmail.com

Zheng(Sandy) Zhang
ZTE Corporation

Email: zhang.zheng@zte.com.cn

Gyan Mishra
Verizon

Email: gyan.s.mishra@verizon.com