

Internet Engineering Task Force (IETF)
Request for Comments: 9624
Category: Standards Track
ISSN: 2070-1721

Z. Zhang
T. Przygienda
Juniper Networks
A. Sajassi
Cisco Systems
J. Rabadan
Nokia
August 2024

EVPN Broadcast, Unknown Unicast, or Multicast (BUM) Using Bit Index Explicit Replication (BIER)

Abstract

This document specifies protocols and procedures for forwarding Broadcast, Unknown Unicast, or Multicast (BUM) traffic of Ethernet VPNs (EVPNs) using Bit Index Explicit Replication (BIER).

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9624>.

Copyright Notice

Copyright (c) 2024 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction
 - 1.1. Terminology
 - 1.2. Requirements Language
2. Use of the PMSI Tunnel Attribute
 - 2.1. IP-Based Tunnel and BIER PHP
 - 2.2. Explicit Tracking
 - 2.2.1. Using IMET/SMET Routes
 - 2.2.2. Using S-PMSI/Leaf A-D Routes
 - 2.3. MPLS Label in the PTA
3. Multihoming Split Horizon
4. Data Plane
 - 4.1. Encapsulation and Transmission

4.1.1.	At a BFIR That Is an Ingress PE
4.1.2.	At a BFIR That Is a P-Tunnel Segmentation Point
4.2.	Disposition
4.2.1.	At a BFER That Is an Egress PE
4.2.2.	At a BFER That Is a P-Tunnel Segmentation Point
5.	IANA Considerations
6.	Security Considerations
7.	References
7.1.	Normative References
7.2.	Informative References
Acknowledgements	
Authors' Addresses	

1. Introduction

[RFC7432] and [RFC8365] specify the protocols and procedures for Ethernet VPNs (EVPNs). For Broadcast, Unknown Unicast, or Multicast (BUM) traffic, provider/underlay tunnels are used to carry the BUM traffic. Several kinds of tunnel technologies can be used as specified in [RFC7432] and [RFC8365], and this document specifies the protocols and procedures to use Bit Index Explicit Replication (BIER) [RFC8279] as provider tunnels for EVPN BUM traffic.

BIER is an architecture that provides optimal multicast forwarding through a "multicast domain" without requiring intermediate routers to maintain any per-flow state or to engage in an explicit tree-building protocol.

The EVPN BUM procedures specified in [RFC7432] and extended in [RFC9572], [RFC9251], and [CMCAST-ENHANCEMENTS] are much aligned with Multicast VPN (MVPN) procedures [RFC6514], and an EVPN Broadcast Domain (BD) corresponds to a VPN in MVPN. As such, this document is also very much aligned with [RFC8556], which specifies MVPN with BIER. For terseness, some background, terms, and concepts are not repeated here. Additionally, some text is borrowed verbatim from [RFC8556].

1.1. Terminology

ES: Ethernet Segment

ESI: Ethernet Segment Identifier

BFR: Bit-Forwarding Router

BFIR: Bit-Forwarding Ingress Router

BFER: Bit-Forwarding Egress Router

BFR-Prefix: An IP address that uniquely identifies a BFR and is routable in a BIER domain.

C-S: A multicast source address identifying a multicast source located at an EVPN customer site. "C-" stands for "Customer-".

C-G: A multicast group address used by an EVPN customer.

C-flow: A customer multicast flow. Each C-flow is identified by the ordered pair (source address, group address), where each address is in the customer's address space. The identifier of a particular C-flow is usually written as (C-S, C-G). Sets of C-flows can be denoted by the use of the "C-*" wildcard (see [RFC6625]), e.g., (C-*, C-G).

P-tunnel: A multicast tunnel through the network of one or more service providers used to transport C-flows. "P-" stands for

"Provider-".

IMET A-D Route: Inclusive Multicast Ethernet Tag Auto-Discovery route. Carried in BGP Update messages, these routes are used to advertise the "default" P-tunnel for a particular BD.

SMET A-D Route: Selective Multicast Ethernet Tag Auto-Discovery route. Carried in BGP Update messages, these routes are used to advertise the C-flows that the advertising Provider Edge (PE) is interested in.

PMSI: Provider Multicast Service Interface [RFC6513]. A conceptual interface used by a PE to send customer multicast traffic to all or some PEs in the same VPN.

I-PMSI: Inclusive PMSI. For all PEs in the same VPN.

S-PMSI: Selective PMSI. For some of the PEs in the same VPN.

I-PMSI A-D Route: Inclusive PMSI Auto-Discovery route used to advertise the tunnels that instantiate an I-PMSI.

S-PMSI A-D Route: Selective PMSI Auto-Discovery route used to advertise that particular C-flows are bound to (i.e., are traveling through) particular P-tunnels.

PTA: PMSI Tunnel Attribute. A BGP attribute used to identify a particular P-tunnel.

VXLAN: Virtual eXtensible Local Area Network [RFC7348]

NVGRE: Network Virtualization Using Generic Routing Encapsulation [RFC7637]

GENEVE: Generic Network Virtualization Encapsulation [RFC8926]

VNI: VXLAN Network Identifier

VSID: Virtual Subnet Identifier

RSVP-TE P2MP: Resource Reservation Protocol for Point-to-Multipoint TE Label Switched Paths (LSPs) [RFC4875]

mLDP P2MP: Multipoint Label Distribution Protocol extensions for Point-to-Multipoint LSPs [RFC6388]

1.2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Use of the PMSI Tunnel Attribute

[RFC7432] specifies that Inclusive Multicast Ethernet Tag (IMET) routes carry a PMSI Tunnel Attribute (PTA) to identify the particular P-tunnel to which one or more BUM flows are being assigned, which is the same as specified in [RFC6514] for MVPN. [RFC8556] specifies the encoding of the PTA for the use of BIER with MVPN. Much of that specification is reused for the use of BIER with EVPN, and much of the text below is borrowed verbatim from [RFC8556].

The PTA contains the following fields:

- * Tunnel Type. The same codepoint 0x0B that IANA has assigned for BIER for MVPN [RFC8556] is used for EVPN as well.
- * Tunnel Identifier. This field contains three subfields for BIER. The text below is exactly as in [RFC8556].
 1. The first subfield is a single octet, containing a BIER sub-domain-id (see [RFC8279]). This indicates that packets sent on the PMSI will be sent on the specified BIER sub-domain. How that sub-domain is chosen is outside the scope of this document.
 2. The second subfield is a two-octet field containing the BFR-id in the sub-domain identified in the first subfield of the router that is constructing the PTA.
 3. The third subfield is the BFR-Prefix (see [RFC8279]) of the router (a BFIR) that is constructing the PTA. The BFR-Prefix will either be a /32 IPv4 address or a /128 IPv6 address. Whether the address is IPv4 or IPv6 can be inferred from the total length of the PTA.

The BFR-Prefix need not be the same IP address that is carried in any other field of the x-PMSI A-D route, even if the BFIR is the originating router of the x-PMSI A-D route.
- * MPLS Label. For EVPN-MPLS [RFC7432], this field contains an upstream-assigned MPLS label. It is assigned by the BFIR. Constraints on how the originating router selects this label are discussed in Section 2.3. For EVPN-VXLAN/NVGRE/GENEVE [RFC8365] [RFC7348] [RFC7637] [RFC8926], this field is a 24-bit VNI/VSID of global significance.
- * Flags. When the tunnel type is BIER, two of the flags in the PTA Flags field are meaningful. Details about the use of these flags can be found in Section 2.2.

- Leaf Info Required per Flow (LIR-pF) [RFC8534]
- Leaf Info Required (LIR)

Note that if a PTA specifying "BIER" is attached to an IMET, S-PMSI A-D, or per-region I-PMSI A-D route, the route MUST NOT be distributed beyond the boundaries of a BIER domain. That is, any routers that receive the route must be in the same BIER domain as the originator of the route. If the originator is in more than one BIER domain, the route must be distributed only within the BIER domain in which the BFR-Prefix in the PTA uniquely identifies the originator. As with all MVPN routes, the distribution of these routes is controlled by the provisioning of Route Targets.

2.1. IP-Based Tunnel and BIER PHP

When VXLAN/NVGRE/GENEVE is used for EVPN, by default, the outer IP header (and UDP header in the case of VXLAN/GENEVE) is not included in the BIER payload, except when it is known a priori that BIER Penultimate Hop Popping (PHP) [BIER-PHP] is used in the BIER domain and the encapsulation (after the BIER header is popped) between the BIER Penultimate Hop and the egress PE does not have a way to indicate the next header is VXLAN/NVGRE/GENEVE. In that case, the full VXLAN/NVGRE/GENEVE encapsulation MUST be used. In the outer IP header, a well-known IP multicast address (224.0.0.122 in the case of IPv4 or FF02:0:0:0:0:0:0:14 in the case of IPv6) is used as the destination address, and the egress PEs MUST be set up to receive and process packets addressed to the destination address. The address is used for all BDs, and the inner VXLAN/NVGRE/GENEVE header will be

used to identify BDs.

2.2. Explicit Tracking

When using BIER to transport an EVPN BUM data packet through a BIER domain, an ingress PE functions as a BFIR (see [RFC8279]). The BFIR must determine the set of BFERs to which the packet needs to be delivered. This can be done in either of two ways as described in the following two sections.

2.2.1. Using IMET/SMET Routes

Both IMET and SMET routes provide explicit tracking functionality.

For an inclusive PMSI, the set of BFERs (egress PEs) includes the originators of all IMET routes for a BD. For a selective PMSI, the set of BFERs (egress PEs) includes the originators of corresponding SMET routes.

The SMET routes do not carry a PTA. When an ingress PE sends traffic on a selective tunnel using BIER, it uses the upstream-assigned label that is advertised in its IMET route.

When only selective forwarding is used for all flows and without tunnel segmentation, SMET routes are used without the need for S-PMSI A-D routes. Otherwise, the procedures in the following section apply.

2.2.2. Using S-PMSI/Leaf A-D Routes

There are two cases where S-PMSI/Leaf A-D routes are used as discussed in the following two sections.

2.2.2.1. Selective Forwarding Only for Some Flows

With the SMET procedure, a PE advertises a SMET route for each (C-S, C-G) or (C-*, C-G) state that it learns on its Attachment Circuits (ACs), and each SMET route is tracked by every PE in the same BD. It may be desired that SMET routes are not used in order to reduce the burden of explicit tracking.

In this case, most multicast traffic will follow the I-PMSI (advertised via the IMET route) and only some flows will follow S-PMSIs. To achieve that, S-PMSI/Leaf A-D routes can be used, as specified in [RFC9572].

The rules specified in Sections 2.2.1 and 2.2.2 of [RFC8556] apply.

2.2.2.2. Tunnel Segmentation

Another case where S-PMSI/Leaf A-D routes are necessary is tunnel segmentation, which is also specified in [RFC9572] and further clarified in [CMCAST-ENHANCEMENTS] for segmentation with SMET routes. This is only applicable to EVPN-MPLS.

The rules specified in Section 2.2.1 of [RFC8556] apply. Section 2.2.2 of [RFC8556] does not apply, because like in MVPN, the LIR-pF flag cannot be used with segmentation.

2.2.2.3. Applicability of Additional MVPN Specifications

As with the MVPN case, "Use of the PMSI Tunnel Attribute in Leaf A-D Routes" (Section 3 of [RFC8556]) applies.

Notice that [RFC8556] refers to procedures specified in [RFC6625] and [RFC8534]. Those two documents were specified for MVPN but apply to

IP multicast payload in EVPN as well.

2.3. MPLS Label in the PTA

Rules in Section 2.1 of [RFC8556] apply, EXCEPT the following three bullets (they do NOT apply to EVPN) in that section:

- * If the two routes do not have the same Address Family Identifier (AFI) value, then their respective PTAs MUST contain different MPLS label values. This ensures that when an egress PE receives a data packet with the given label, the egress PE can infer from the label whether the payload is an IPv4 packet or an IPv6 packet.
- * If the BFIR is an ingress PE supporting MVPN extranet [RFC7900] functionality, and if the two routes originate from different VRFs on this ingress PE, then the respective PTAs of the two routes MUST contain different MPLS label values.
- * If the BFIR is an ingress PE supporting the "Extranet Separation" feature of MVPN extranet (see Section 7.3 of [RFC7900]), and if one of the routes carries the "Extranet Separation" extended community but the other does not, then the respective PTAs of the two routes MUST contain different MPLS label values.

3. Multihoming Split Horizon

For EVPN-MPLS, [RFC7432] specifies the use of ESI labels to identify the ES from which a BUM packet originates. A PE receiving that packet from the core side will not forward it to the same ES. The procedure works for both Ingress Replication (IR) and RSVP-TE/mLDP P2MP tunnels, using downstream- and upstream-assigned ESI labels, respectively. For EVPN-VXLAN/NVGRE/GENEVE, [RFC8365] specifies local bias procedures, where a PE receiving a BUM packet from the core side knows the ingress PE due to encapsulation; therefore, the PE does not forward the packet to any multihoming ESes that the ingress PE is on. This is because the ingress PE already forwarded the packet to those ESes, regardless of whether the ingress PE is a Designated Forwarder for those ESes.

With BIER, the local bias procedure still applies for EVPN-VXLAN/NVGRE/GENEVE, as the BFIR-id in the BIER header identifies the ingress PE. For EVPN-MPLS, ESI label procedures also still apply, though two upstream-assigned labels will be used (one for identifying the BD and one for identifying the ES) -- the same as in the case of using a single P2MP tunnel for multiple BDs. The BFIR-id in the BIER header identifies the ingress PE that assigned those two labels.

4. Data Plane

Like MVPN, the EVPN application plays the role of the "multicast flow overlay" as described in [RFC8279].

4.1. Encapsulation and Transmission

A BFIR could be either an ingress PE or a P-tunnel segmentation point. The procedures are slightly different as described below.

4.1.1. At a BFIR That Is an Ingress PE

To transmit a BUM data packet, an ingress PE first determines the route matched for transmission and routes for tracking leaves according to the following rules.

1. If selective forwarding is not used or is not an IP multicast packet after the Ethernet header, the IMET route originated for the BD by the ingress PE is the route matched for transmission.

Leaf-tracking routes are all other received IMET routes for the BD.

2. Otherwise, if selective forwarding is used for all IP multicast traffic based on SMET routes, the IMET route originated for the BD by the ingress PE is the route matched for transmission. Received SMET routes for the BD, whose source and destination address fields match the packet's source and destination IP address, are leaf-tracking routes.
3. Otherwise, the route matched for transmission is the S-PMSI A-D route originated by the ingress PE for the BD, whose source and destination address fields match the packet's source and destination IP address and have a PTA specifying a valid tunnel type that is not "no tunnel info". Leaf-tracking routes are determined as follows:
 - a. If the match for the transmission route carries a PTA that has the LIR flag set but does not have the LIR-pF flag set, the routes matched for tracking are Leaf A-D routes whose Route Key field is identical to the NLRI of the S-PMSI A-D route.
 - b. If the match for the transmission route carries a PTA that has the LIR-pF flag, the leaf-tracking routes are Leaf A-D routes whose Route Key field is derived from the NLRI of the S-PMSI A-D route according to the procedures described in Section 5.2 of [RFC8534].

Note that in both cases, SMET routes may be used in lieu of Leaf A-D routes, as a PE may omit the Leaf A-D route in response to an S-PMSI A-D route with the LIR or LIR-pF bit set if a SMET route with the corresponding Tag, Source, and Group fields is already originated [RFC9572]. In particular, in the second case above, even though the SMET route does not have a PTA attached, it is still considered a Leaf A-D route in response to a wildcard S-PMSI A-D route with the LIR-pF bit set.

4. Otherwise, the route matched for transmission and leaf-tracking routes are determined as in rule 1.

If no route is matched for transmission, the packet is not forwarded onto a P-tunnel. If the tunnel that the ingress determines to use based on the route matched for transmission (and considering interworking with PEs that do not support certain tunnel types per procedures in [RFC9251]) requires leaf tracking (e.g., Ingress Replication, RSVP-TE P2MP tunnel, or BIER) but there are no leaf-tracking routes, the packet will not be forwarded onto a P-tunnel either.

The following text assumes that BIER is the determined tunnel type. The ingress PE pushes an upstream-assigned ESI label per [RFC7432] if the following conditions are all met:

- * The packet is received on a multihomed ES.
- * It is EVPN-MPLS.
- * The ESI label procedure is used for split horizon.

The MPLS label from the PTA of the route matched for transmission is then pushed onto the packet's label stack for EVPN-MPLS. For EVPN-VXLAN/NVGRE/GENEVE, a VXLAN/NVGRE/GENEVE header is prepended to the packet with the VNI/VSID set to the value in the PTA's Label field, and then an IP/UDP header is prepended if needed (e.g., for PHP purposes).

Then, the packet is encapsulated in a BIER header and forwarded according to the procedures of [RFC8279] and [RFC8296]. Specifically, see "Imposing and Processing the BIER Encapsulation" (Section 3 of [RFC8296]). The Proto field in the BIER header is set to 2 in the case of EVPN-MPLS, 7/8/9 in the case of EVPN-VXLAN/NVGRE/GENEVE (Section 5) when an IP header is not used, or 4/6 if an IP header is used for EVPN-VXLAN/NVGRE/GENEVE.

To create the proper BIER header for a given packet, the BFIR must know all the BFERs that need to receive that packet. This is determined from the set of leaf-tracking routes.

4.1.2. At a BFIR That Is a P-Tunnel Segmentation Point

In this case, the encapsulation for the upstream segment of the P-tunnel includes (among other things) a label that identifies the x-PMSI or IMET A-D route that is the match for reception on the upstream segment. The segmentation point re-advertised the route into one or more downstream regions. Each instance of the re-advertised route for a downstream region has a PTA that specifies the tunnel for that region. For any particular downstream region, the route matched for transmission is the re-advertised route, and the leaf-tracking routes are determined as follows, if needed, for the tunnel type:

- * If the route matched for transmission is an x-PMSI route, it must have the LIR flag set in its PTA, and the leaf-tracking routes are all the matching Leaf A-D and SMET routes received in the downstream region.
- * If the route matched for transmission is an IMET route, the leaf-tracking routes are all the IMET routes for the same BD received in the downstream region.

If the downstream region uses BIER, the packet is forwarded as follows: the upstream segmentation's encapsulation is removed and the above-mentioned label is swapped to the upstream-assigned label in the PTA of the route matched for transmission, and then a BIER header is imposed as in Section 4.1.1.

4.2. Disposition

The same procedures in Section 4.2 of [RFC8556] are followed for EVPN-MPLS, except for some EVPN specifics that are discussed in the following two subsections of this document.

For EVPN-VXLAN/NVGRE/GENEVE, the only differences are that the payload is VXLAN/NVGRE/GENEVE (with or without an IP header) and the VNI/VSID field in the VXLAN/NVGRE/GENEVE header is used to determine the corresponding BD.

4.2.1. At a BFER That Is an Egress PE

Once the corresponding BD is determined from the upstream-assigned label or VNI/VSID, EVPN forwarding procedures per [RFC7432] or [RFC8365] are followed. In the case of EVPN-MPLS, if there is an inner label in the label stack following the BIER header, that inner label is considered the upstream-assigned ESI label for split-horizon purposes.

4.2.2. At a BFER That Is a P-Tunnel Segmentation Point

This is only applicable to EVPN-MPLS. The same procedures in Section 4.2.2 of [RFC8556] are followed, subject to multihoming procedures specified in [RFC9572].

5. IANA Considerations

Per this document, IANA has registered the following three values in the "BIER Next Protocol Identifiers" registry:

Value	Description	Reference
7	Payload is VXLAN encapsulated (no IP/UDP header)	RFC 9624
8	Payload is NVGRE encapsulated (no IP header)	RFC 9624
9	Payload is GENEVE encapsulated (no IP/UDP header)	RFC 9624

Table 1: BIER Next Protocol Identifiers Registry

IANA has also assigned an IPv4 and an IPv6 multicast address for the case discussed in Section 2.1.

The following entry has been added to the "Local Network Control Block (224.0.0.0 - 224.0.0.255 (224.0.0/24))" registry for IPv4:

Address(es): 224.0.0.122
Description: Network Virtualization Overlay (NVO) BUM Traffic
Reference: RFC 9624

The following entry has been added to the "Link-Local Scope Multicast Addresses" registry for IPv6:

Address(es): FF02:0:0:0:0:0:0:14
Description: Network Virtualization Overlay (NVO) BUM Traffic
Reference: RFC 9624

6. Security Considerations

This document is about using BIER as provider tunnels for EVPN. It is very similar to using BIER as MVPN provider tunnels and does not introduce additional security implications beyond what have been discussed in the EVPN base protocol specification [RFC7432] and MVPN using BIER [RFC8556].

7. References

7.1. Normative References

- [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>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.
- [RFC6625] Rosen, E., Ed., Rekhter, Y., Ed., Hendrickx, W., and R. Qiu, "Wildcards in Multicast VPN Auto-Discovery Routes",

RFC 6625, DOI 10.17487/RFC6625, May 2012,
<<https://www.rfc-editor.org/info/rfc6625>>.

- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC7900] Rekhter, Y., Ed., Rosen, E., Ed., Aggarwal, R., Cai, Y., and T. Morin, "Extranet Multicast in BGP/IP MPLS VPNs", RFC 7900, DOI 10.17487/RFC7900, June 2016, <<https://www.rfc-editor.org/info/rfc7900>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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>>.
- [RFC8365] Sajassi, A., Ed., Drake, J., Ed., Bitar, N., Shekhar, R., Uttaro, J., and W. Henderickx, "A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)", RFC 8365, DOI 10.17487/RFC8365, March 2018, <<https://www.rfc-editor.org/info/rfc8365>>.
- [RFC8534] Dolganow, A., Kotalwar, J., Rosen, E., Ed., and Z. Zhang, "Explicit Tracking with Wildcard Routes in Multicast VPN", RFC 8534, DOI 10.17487/RFC8534, February 2019, <<https://www.rfc-editor.org/info/rfc8534>>.
- [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>>.
- [RFC8926] Gross, J., Ed., Ganga, I., Ed., and T. Sridhar, Ed., "Geneve: Generic Network Virtualization Encapsulation", RFC 8926, DOI 10.17487/RFC8926, November 2020, <<https://www.rfc-editor.org/info/rfc8926>>.
- [RFC9251] Sajassi, A., Thoria, S., Mishra, M., Patel, K., Drake, J., and W. Lin, "Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Proxies for Ethernet VPN (EVPN)", RFC 9251, DOI 10.17487/RFC9251, June 2022, <<https://www.rfc-editor.org/info/rfc9251>>.
- [RFC9572] Zhang, Z., Lin, W., Rabadan, J., Patel, K., and A. Sajassi, "Updates to EVPN Broadcast, Unknown Unicast, or Multicast (BUM) Procedures", RFC 9572, DOI 10.17487/RFC9572, May 2024, <<https://www.rfc-editor.org/info/rfc9572>>.

7.2. Informative References

- [BIER-PHP] Zhang, Z., "BIER Penultimate Hop Popping", Work in Progress, Internet-Draft, draft-ietf-bier-php-11, 6

February 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-php-11>>.

[CMCAST-ENHANCEMENTS]

Zhang, Z., Kebler, R., Lin, W., and E. Rosen, "MVPN/EVPN C-Multicast Routes Enhancements", Work in Progress, Internet-Draft, draft-zzhang-bess-mvpn-evpn-cmcast-enhancements-04, 17 March 2024, <<https://datatracker.ietf.org/doc/html/draft-zzhang-bess-mvpn-evpn-cmcast-enhancements-04>>.

[RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <<https://www.rfc-editor.org/info/rfc4875>>.

[RFC6388] Wijnands, IJ., Ed., Minei, I., Ed., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, DOI 10.17487/RFC6388, November 2011, <<https://www.rfc-editor.org/info/rfc6388>>.

[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, DOI 10.17487/RFC7348, August 2014, <<https://www.rfc-editor.org/info/rfc7348>>.

[RFC7637] Garg, P., Ed. and Y. Wang, Ed., "NVGRE: Network Virtualization Using Generic Routing Encapsulation", RFC 7637, DOI 10.17487/RFC7637, September 2015, <<https://www.rfc-editor.org/info/rfc7637>>.

Acknowledgements

The authors thank Eric Rosen for his review and suggestions. Additionally, much of the text is borrowed verbatim from [RFC8556].

Authors' Addresses

Zhaohui Zhang
Juniper Networks
Email: zzhang@juniper.net

Tony Przygienda
Juniper Networks
Email: prz@juniper.net

Ali Sajassi
Cisco Systems
Email: sajassi@cisco.com

Jorge Rabadan
Nokia
Email: jorge.rabadan@nokia.com