

BESS Working Group
Internet-Draft
Obsoletes: 8317 (if approved)
Updates: 7385 (if approved)
Intended status: Standards Track
Expires: 3 January 2026

A. Sajassi, Ed.
Cisco
J. Rabadan
Nokia
J. Drake
Independent
A. Appachi gounder
Google
A. Bamberger
Arista Networks
2 July 2025

Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) and Provider
Backbone Bridging EVPN (PBB-EVPN)
draft-sajassi-bess-rfc8317bis-04

Abstract

The MEF Forum (MEF) has defined a rooted-multipoint Ethernet service known as Ethernet-Tree (E-Tree). A solution framework for supporting this service in MPLS networks is described in RFC7387, "A Framework for Ethernet-Tree (E-Tree) Service over a Multiprotocol Label Switching (MPLS) Network". This document discusses how those functional requirements can be met with a solution based on RFC7432, "BGP MPLS Based Ethernet VPN (EVPN)", with some extensions and a description of how such a solution can offer a more efficient implementation of these functions than that of RFC7796, "Ethernet-Tree (E-Tree) Support in Virtual Private LAN Service (VPLS)". This document makes use of the most significant bit of the Tunnel Type field (in the P-Multicast Service Interface (PMSI) Tunnel attribute) governed by the IANA registry created by RFC7385; hence, it updates RFC7385 accordingly. This document obsoletes RFC8317.

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.

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 3 January 2026.

Copyright Notice

Copyright (c) 2025 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	3
2. Requirements Language	4
3. Terminology	4
4. E-Tree Scenarios	6
4.1. Scenario 1: Leaf or Root Site(s) per PE	7
4.2. Scenario 2: Leaf or Root Site(s) per AC	8
4.3. Scenario 3: Leaf or Root Site(s) per MAC Address	9
5. Operation for EVPN	10
5.1. Known Unicast Traffic	10
5.2. BUM Traffic with MPLS Encapsulation	11
5.2.1. BUM Traffic Originated from a Single-Homed Site on a Leaf AC	12
5.2.2. BUM Traffic Originated from a Single-Homed Site on a Root AC	12
5.2.3. BUM Traffic Originated from a Multihomed Site on a Leaf AC	13
5.2.4. BUM Traffic Originated from a Multihomed Site on a Root AC	13
5.3. BUM Traffic with VxLAN or GENEVE Encapsulation	13
5.4. BUM Traffic with SRv6 Encapsulation	14

5.5.	E-Tree Traffic Flows for EVPN	15
5.5.1.	E-Tree with MAC Learning	15
5.5.2.	E-Tree without MAC Learning	16
5.6.	Adaptive Ingress and Egress Filtering of BUM Traffic	17
5.6.1.	Control Plane Operation: Advertising PE	17
5.6.2.	Control Plane Operation: Receiving PE	18
5.6.3.	Data Plane Operation	19
6.	Operation for PBB-EVPN	20
6.1.	Known Unicast Traffic	21
6.2.	BUM Traffic	21
6.3.	E-Tree without MAC Learning	22
7.	BGP Encoding	22
7.1.	E-Tree Extended Community	22
7.2.	PMSI Tunnel Attribute	23
8.	Security Considerations	25
9.	IANA Considerations	25
9.1.	Considerations for PMSI Tunnel Types	25
10.	Additional Authors from RFC8317	26
	Acknowledgements	26
	References	27
	Normative References	27
	Informative References	28
	Authors' Addresses	29

1. Introduction

The MEF Forum (MEF) has defined a rooted-multipoint Ethernet service known as Ethernet-Tree (E-Tree) [MEF6.1]. In an E-Tree service, a customer site that is typically represented by an Attachment Circuit (AC) (e.g., an 802.1Q VLAN tag), is labeled as either a Root or a Leaf site. A customer site may also be represented by a Media Access Control (MAC) address along with a VLAN tag. Root sites can communicate with all other customer sites (both Root and Leaf sites). However, Leaf sites can communicate with Root sites but not with other Leaf sites. In this document, unless explicitly mentioned otherwise, a site is always represented by an AC.

[RFC7387] describes a solution framework for supporting E-Tree service in MPLS networks. This document identifies the functional components of an overall solution to emulate E-Tree services in MPLS networks and supplements the multipoint-to-multipoint Ethernet LAN (E-LAN) services specified in [RFC7432] and [RFC7623].

[RFC7432] defines EVPN, a solution for multipoint Layer 2 Virtual Private Network (L2VPN) services with advanced multihoming capabilities that uses BGP for distributing customer/client MAC address reachability information over the MPLS/IP network. [RFC7623] combines the functionality of EVPN with IEEE 802.1ah Provider Backbone Bridging (PBB) for MAC address scalability.

This document discusses how the functional requirements for E-Tree service can be met with a solution based on EVPN [RFC7432] and PBB-EVPN [RFC7623] with some extensions to their procedures and BGP attributes. Such a solution based on PBB-EVPN or EVPN can offer a more efficient implementation of these functions than that of [RFC7796], "Ethernet-Tree (E-Tree) Support in Virtual Private LAN Service (VPLS)". This efficiency is achieved by performing filtering of unicast traffic at the ingress Provider Edge (PE) nodes as opposed to egress filtering where the traffic is sent through the network and gets filtered and discarded at the egress PE nodes. The details of this ingress filtering are described in Section 5.1. Since this document specifies a solution based on [RFC7432], the knowledge of that document is a prerequisite. This document makes use of the most significant bit of the Tunnel Type field (in the PMSI Tunnel attribute) governed by the IANA registry created by [RFC7385]; hence, it updates [RFC7385] accordingly. Section 4 discusses E-Tree scenarios, Section 5 and Section 6 describe E-Tree solutions for EVPN and PBB-EVPN (respectively), and Section 7 covers BGP encoding for E-Tree solutions.

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.

3. Terminology

BD: Broadcast Domain. In a bridged network, the broadcast domain corresponds to a Virtual LAN (VLAN). In this document, "BD", "subnet", and VLAN are equivalent terms, and wherever "subnet" is used, it means "IP subnet". As per [RFC7432], an EVI consists of a single BD or multiple BDs. In the case of VLAN-bundle and VLAN-based service models, a BD is equivalent to an EVI. In the case of a VLAN-aware bundle service model, an EVI contains multiple BDs.

EVI: EVPN Instance spanning NVE/PE devices that are participating in

that EVPN. As per [RFC7432], an EVI consists of a single BD or multiple BDs. In the case of VLAN-bundle and VLAN-based service models (see [RFC7432]), an EVI is equivalent to a BD. In the case of a VLAN-aware bundle service model, an EVI contains multiple BDs.

MAC-VRF: A Virtual Routing and Forwarding table for Media Access Control (MAC) addresses on a PE.

Ethernet Segment (ES): When a customer site (device or network) is connected to one or more PEs via a set of Ethernet links, then that set of links is referred to as an 'Ethernet segment'.

Ethernet Segment Identifier (ESI): A unique non-zero identifier that identifies an Ethernet segment is called an 'Ethernet Segment Identifier'.

VID: VLAN Identifier.

Ethernet Tag: Used to represent a BD that is configured on a given ES for the purposes of DF election and <BD, BD> identification for frames received from the CE. Note that any of the following may be used to represent a BD: VIDs (including Q-in-Q tags), configured IDs, VNIs (Virtual Extensible Local Area Network (VXLAN) Network Identifiers), normalized VIDs, I-SIDs (Service Instance Identifiers), etc., as long as the representation of the BDs is configured consistently across the multihomed PEs attached to that ES.

Ethernet Tag ID: Normalized network wide ID that is used to identify a BD within a BD and carried in EVPN routes.

MP2MP: Multipoint to Multipoint.

MP2P: Multipoint to Point.

P2MP: Point to Multipoint.

P2P: Point to Point.

PE: Provider Edge device.

Single-Active Redundancy Mode: When only a single PE, among all the PEs attached to an Ethernet segment, is allowed to forward traffic to/from that Ethernet segment for a given VLAN, then the Ethernet segment is defined to be operating in Single-Active redundancy mode.

All-Active Redundancy Mode: When all PEs attached to an Ethernet segment are allowed to forward known unicast traffic to/from that Ethernet segment for a given VLAN, then the Ethernet segment is defined to be operating in All-Active redundancy mode.

BUM: Broadcast, unknown unicast, and multicast.

DF: Designated Forwarder

Backup-DF (BDF): Backup-Designated Forwarder.

Non-DF (NDF): Non-Designated Forwarder.

AC: Attachment Circuit

NVO: Network Virtualization Overlay as described in [RFC8365]

IRB: Integrated Routing and Bridging interface, with EVPN procedures described in [RFC9135]

IMET route: Inclusive Multicast Ethernet Tag route" described in [RFC7432]

DCB label: Domain-wide Common Block label

4. E-Tree Scenarios

This document categorizes E-Tree scenarios into the following three categories, depending on the nature of the Root/Leaf site association:

- * Scenario 1: either Leaf or Root site(s) per PE
- * Scenario 2: either Leaf or Root site(s) per Attachment Circuit (AC)
- * Scenario 3: either Leaf or Root site(s) per MAC address

Scenarios 1 and 2 are of utmost importance and are covered extensively in this document. Since publication of [RFC8317], no network application for scenario 3 has been identified because E-Tree segmentation for scenario 3 can be enforced for known unicast traffic, but it cannot be enforced for BUM traffic. This limits the applicability of scenario-3 significantly for EVPN services.

4.1. Scenario 1: Leaf or Root Site(s) per PE

In this scenario, a PE may receive traffic from either Root ACs or Leaf ACs for a given Broadcast Domain (BD)(i.e., EVI or EVI+VLAN) but not both. In other words, a given BD on a Provider Edge (PE) device is either associated with Root(s) or Leaf(s). The PE may have both Root and Leaf ACs, albeit for different BDs.

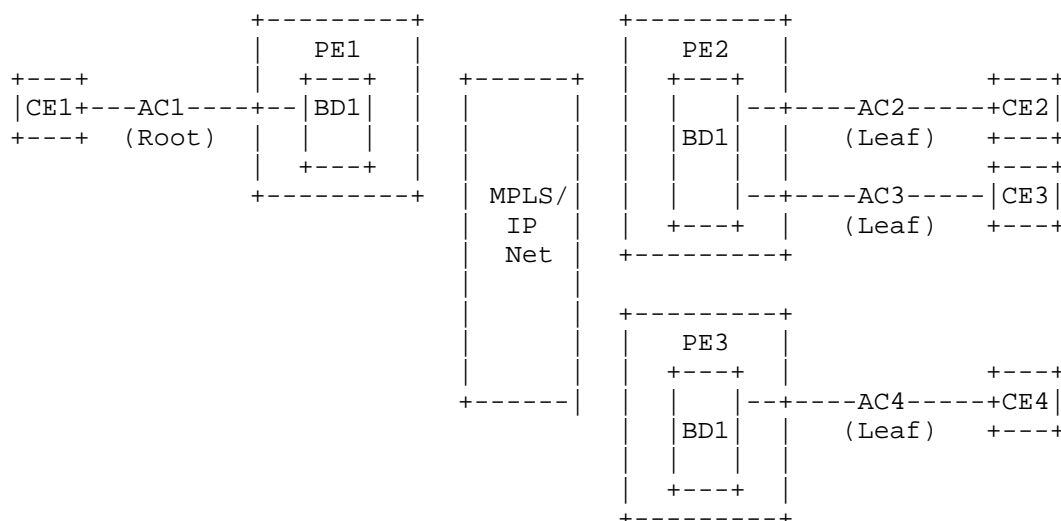


Figure 1: Scenario 1

In this scenario, even though there is no EVPN data-plane communications among leaf PEs belonging to the same BD, there needs to be EVPN control-plane communications among these PEs because of host mobility (i.e., a host can move from one leaf PE to another leaf PE in the same BD). Therefore, EVPN unicast MAC/IP routes need to be exchanged among the leaf PEs to allow for EVPN MAC mobility procedures to get executed properly. This implies a single Route Target just like baseline EVPN per EVI must be used for this scenario.

Furthermore, in this scenario, known unicast traffic from a leaf PE destined to another leaf PE must be dropped at the ingress PE (i.e., known unicast traffic from a leaf PE is only allowed to a root PE). This ingress filtering of known unicast traffic is achieved by signaling extensions to EVPN for coloring EVPN MAC/IP routes as described in Section 5. Because of simple E-Tree topology for this scenario where a given PE in a BD is either root or leaf (but not both), the ingress filtering can be extended for BUM traffic in case of ingress replication. This ingress filtering of BUM traffic can be

In this scenario, ingress filtering for known unicast traffic is performed just like scenario 1. However, ingress-only filtering for BUM traffic is not possible for this scenario because a participant PE (e.g., PE2 in the above figure) can have both leaf and root ACs and thus need to receive the BUM traffic and perform egress filtering.

In order to perform egress filtering for BUM traffic received at the egress PE, the ingress PE needs to color the BUM traffic in data-plane to indicate if the traffic is coming from a Root or a Leaf. The egress PE uses this indication in data-plane in its egress filtering decision as described in Section 5.2.

In this scenario, the transmission of BUM traffic to egress PEs (in a given BD) that are only configured with leaf ACs, can be optimized by ingress filtering of BUM traffic to those PEs. However, because of dynamic nature of AC leaf/root activation on a PE, such ingress filtering requires extensions to EVPN signaling as described in Section 5.6. This adaptive ingress filtering optimization for BUM traffic is optional and it is only applicable to ingress replication tunnels. For a P2MP tunnel sourced from a leaf PE, other leaf PEs in that BD should simply avoid joining that tree.

4.3. Scenario 3: Leaf or Root Site(s) per MAC Address

In this scenario, a customer Root or Leaf site is represented by a MAC address on an AC and a PE may receive traffic from both Root and Leaf sites on that AC for a BD. This scenario is not covered in either [RFC7387] or [MEF6.1]; however, it is covered in this document for the sake of completeness. In this scenario, since an AC carries traffic from both Root and Leaf sites, the granularity at which Root or Leaf sites are identified is on a per-MAC-address basis. This scenario is considered in this document for EVPN service with only known unicast traffic because the Designated Forwarder (DF) filtering per [RFC7432] would not be compatible with the required egress filtering; that is, Broadcast, Unknown Unicast, and Multicast (BUM) traffic is not supported in this scenario; it is dropped by the ingress PE.

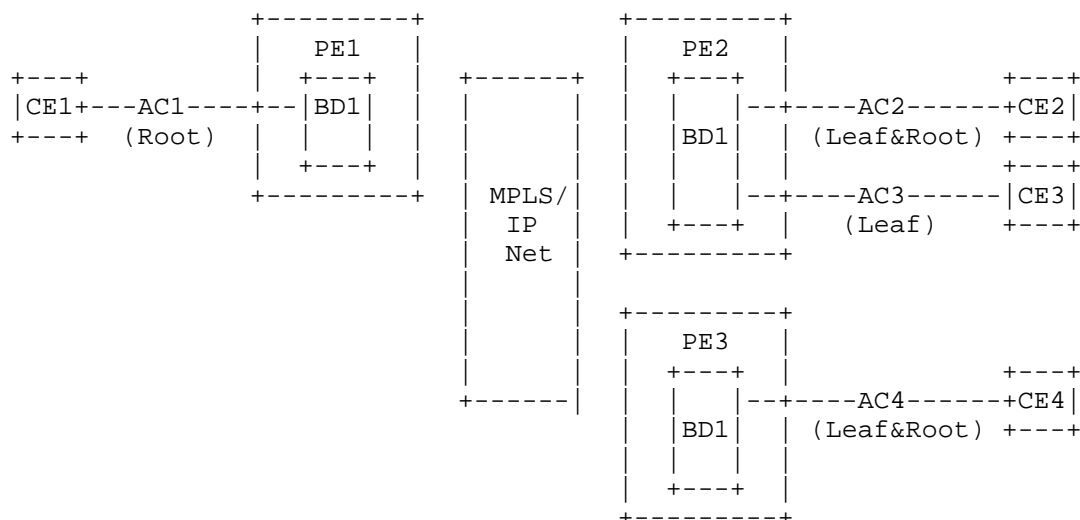


Figure 3: Scenario 3

5. Operation for EVPN

[RFC7432] defines the notion of the Ethernet Segment Identifier (ESI) MPLS label used for split-horizon filtering of BUM traffic at the egress PE. Such egress filtering capabilities can be leveraged in provision of E-Tree services, as it will be described in Section 5.2 for BUM traffic when MPLS encapsulation is used. For known unicast traffic, additional extensions to [RFC7432] are needed (i.e., a new BGP extended community for Leaf-Indication described in Section 7) in order to enable ingress filtering as described in detail in the following sections.

5.1. Known Unicast Traffic

In EVPN, MAC learning is performed in the control plane via advertisement of BGP routes. Because of this, the filtering needed by an E-Tree service for known unicast traffic can be performed at the ingress PE, thus providing very efficient filtering and avoiding sending known unicast traffic over the MPLS/IP core to be filtered at the egress PE, as is done in traditional E-Tree solutions (i.e., E-Tree for VPLS [RFC7796]).

To provide such ingress filtering for known unicast traffic, a PE MUST indicate to other PEs what kind of sites (Root or Leaf) its MAC addresses are associated with. This is done by advertising a Leaf-Indication flag via a new E-Tree extended community specified in Section 7 along with each of its MAC/IP Advertisement routes learned

from a Leaf site. This new extended community MUST be advertised with MAC/IP Advertisement routes learned from a Leaf site. The lack of such a flag indicates that the MAC address is associated with a Root site. This scheme applies to all scenarios described in Section 4.

Tagging MAC addresses with a Leaf-Indication enables remote PEs to perform ingress filtering for known unicast traffic; that is, on the ingress PE, the MAC destination address lookup yields (in addition to the forwarding adjacency) a flag that indicates whether or not the target MAC is associated with a Leaf site. The ingress PE cross-checks this flag with the status of the originating AC, and if both are Leafs, then the packet is not forwarded.

In a situation where MAC moves are allowed among Leaf and Root sites (e.g., non-static MAC), PEs can receive multiple MAC/IP Advertisement routes for the same MAC address with different Root or Leaf-Indications (and possibly different ESIs for multihoming scenarios). In such situations, MAC mobility procedures (see Section 15 of [RFC7432]) take precedence to first identify the location of the MAC before associating that MAC with a Root or a Leaf site.

5.2. BUM Traffic with MPLS Encapsulation

This section specifies the procedure for egress filtering of BUM traffic with MPLS encapsulation. To support scenario-2 efficiently, egress filtering of BUM traffic is required as described below. In order to apply the proper egress filtering, which varies based on whether a packet is sent from a Leaf AC or a Root AC, the MPLS-encapsulated frames MUST be tagged with an indication of when they originated from a Leaf AC (i.e., to be tagged with a Leaf label as specified in Section 7). This Leaf label allows for disposition PE (e.g., egress PE) to perform the necessary egress filtering function in a data plane similar to the ESI label in [RFC7432]. The allocation of the Leaf label can be domain wide (i.e., DCB labels - [I-D.ietf-bess-mvpn-evpn-aggregation-label]) or it can be on a per-PE basis (e.g., independent of ESI and EVI) as described in the following sections.

The Leaf label can be upstream assigned for Point-to-Multipoint (P2MP) Label Switched Path (LSP) or downstream assigned for Ingress Replication tunnels. The main difference between a downstream- and upstream-assigned Leaf label is that, in the case of downstream-assigned Leaf labels, not all egress PE devices need to receive the label in MPLS-encapsulated BUM packets, just like the ESI label for Ingress Replication procedures defined in [RFC7432].

On the ingress PE, the PE needs to place all its Leaf ACs for a given bridge domain in a single split-horizon group in order to prevent intra-PE forwarding of BUM traffic among its Leaf ACs.

There are four scenarios to consider as follows. In all these scenarios, the ingress PE imposes the right MPLS label associated with the originated Ethernet Segment (ES) depending on whether the Ethernet frame originated from a Root or a Leaf site on that Ethernet Segment (ESI label or Leaf label). The mechanism by which the PE identifies whether a given frame originated from a Root or a Leaf site on the segment is based on the AC identifier for that segment (e.g., Ethernet Tag of the frame for 802.1Q frames). Other mechanisms for identifying Root or Leaf sites, such as the use of the source MAC address of the receiving frame, are optional. The scenarios below are described in context of a Root/Leaf AC, however, they can be extended to the Root/Leaf MAC address if needed.

5.2.1. BUM Traffic Originated from a Single-Homed Site on a Leaf AC

In this scenario, the ingress PE adds a Leaf label advertised using the E-Tree extended community (see Section 7), which indicates a Leaf site. This Leaf label, used for single-homing scenarios, is not on a per-ES basis but rather on a per PE basis (i.e., a single Leaf MPLS label is used for all single-homed ESs on that PE). This Leaf label is advertised to other PE devices using the E-Tree extended community (see Section 7) along with an Ethernet Auto-Discovery per ES (EAD-ES) route with an ESI of zero and a set of RTs corresponding to all BDs on the PE where each BD has at least one Leaf site. Multiple EAD-ES routes will need to be advertised if the number of RTs that need to be carried exceed the limit on a single route per [RFC7432]. The ESI for the EAD-ES route is set to zero to indicate single-homed sites.

When a PE receives this special Leaf label in the data path, it blocks the packet if the destination AC is of type Leaf; otherwise, it forwards the packet.

5.2.2. BUM Traffic Originated from a Single-Homed Site on a Root AC

In this scenario, the ingress PE does not add any ESI or Leaf labels and it operates per the procedures in [RFC7432].

5.2.3. BUM Traffic Originated from a Multihomed Site on a Leaf AC

In this scenario, it is assumed that while different ACs (VLANs) on the same ES could have a different Root/Leaf designation (some being Roots and some being Leafs), the same VLAN does have the same Root/Leaf designation on all PEs on the same ES. Furthermore, it is assumed that there is no forwarding among subnets (i.e., the service is EVPN L2 and not EVPN Integrated Routing and Bridging (IRB) [RFC9135]). IRB use cases described in [RFC9135] are outside the scope of this document.

In this scenario, if a multicast or broadcast packet is originated from a Leaf AC, then it only needs to carry a Leaf label as described in Section 5.2.1. This label is sufficient in providing the necessary egress filtering of BUM traffic from getting sent to Leaf ACs, including the Leaf AC on the same ES.

5.2.4. BUM Traffic Originated from a Multihomed Site on a Root AC

In this scenario, both the ingress and egress PE devices follow the procedure defined in [RFC7432] for adding and/or processing an ESI MPLS label; that is, existing procedures for BUM traffic in [RFC7432] are sufficient and there is no need to add a Leaf label.

5.3. BUM Traffic with VxLAN or GENEVE Encapsulation

This section specifies the procedure for egress filtering of BUM traffic with overlay encapsulation for IP tunnels such as VxLAN or GENEVE. As mentioned previously, in order to support scenario-2 efficiently, egress filtering of BUM traffic is required, and in order to support egress filtering, coloring of BUM traffic in data-plane is required to indicate whether the source of the traffic is from a leaf site or a root site.

In order to have a uniform coloring mechanism across all non-MPLS overlay encapsulation types such as VxLAN, GPE, or GENEVE, this specification proposes the use of VNI as the primary mechanism for such coloring of BUM traffic similar to the use of MPLS label when MPLS encapsulation is used. A PE that is configured with a leaf AC for a given BD, advertises its IMET route for that BD along with Tunnel Encapsulation EC indicating IP-based tunnel type (e.g., VxLAN or GENEVE) and E-Tree extended community with leaf-indication flag set and a valid VNI (not zero and not 0xFFFFFFFF). The leaf VNI is encoded in E-Tree extended community and it is in addition to the base VNI which is sent along with EVPN IMET route in PMSI Tunnel attribute. The leaf VNI is often domain-wide VNI; however, if needed it can be downstream assigned. When the receiving PE receives an IMET route with the Tunnel Encapsulation EC and E-Tree EC, it checks

to see if the leaf-indication flag is set. If it is set, then it checks the received leaf VNI against its locally configured leaf VNI for that BD (for domain-wide VNI assignment). If there is no match, the IMET route is discarded and an error message is logged. The interpretation of Leaf VNI/Label field of E-Tree EC is based on the Tunnel Encapsulation EC received with the IMET route. An IP-based tunnel type such as VxLAN or GENEVE, indicates to the receiving PE that the Leaf VNI/Label field to be interpreted as a 24-bit Leaf VNI.

The imposition PE, when wants to send BUM traffic, it uses the leaf VNI if the traffic is sourced from a leaf site. If the BUM traffic is sourced from a root site, then existing base VNI is used. The leaf VNI identifies both the BD and the leaf role. The disposition PE, when receives a VxLAN or GENEVE encapsulated packet with leaf VNI, performs egress filtering accordingly - i.e., it drops the packet at the egress leaf ACs and passes it at the egress root ACs.

Optionally, coloring of BUM traffic with leaf indication in data-plane MAY be done via GPE [I-D.ietf-nvo3-vxlan-gpe] or GENEVE [RFC8926] header by using a single bit from the reserved field in those headers. To signal such coloring mechanism, a PE advertises its IMET route for a given BD along with E-Tree EC with leaf-indication flag set and the VNI of 0xFFFFFFFF.

5.4. BUM Traffic with SRv6 Encapsulation

This section specifies the procedure for egress filtering of BUM traffic with SRv6 encapsulation. The procedure is analogous to the one described in Section 5.2 for MPLS tunnels and ingress replication, only replacing the Leaf labels in MPLS-encapsulated frames with SRv6 encapsulated frames with ultimate destination being the SRv6 SID of End.DT2M behavior [RFC8986] along with the Arg.FE2 arguments encoded in it.

The Arg.FE2 argument is allocated by the egress PE and, as in Section 5.2, signaled in Ethernet A-D per ES routes with ESI of zero and a set of route targets corresponding to all BDs on the PE where each BD has at least one Leaf site.

In the cases where the ingress PE needs to send BUM traffic to the egress PE that advertised the Arg.FE2 argument, and that traffic is coming from a Leaf Attachment Circuit, the ingress PE encodes the Arg.FE2 value as part of the SID value that is used as the ultimate destination of the SRv6 encapsulated BUM frame. Those are the same cases as in Section 5.2 for traffic originated from a Leaf AC (single or multi-homed).

The encoding of the Arg.FE2 argument in Ethernet A-D per ES routes, the encoding of the End.DT2M SID without the ARG part in the IMET routes, as well as the processing of the routes on the ingress PE and error handling on reception are specified in [I-D.ietf-bess-bgp-srv6-args].

5.5. E-Tree Traffic Flows for EVPN

Per [RFC7387], a generic E-Tree service supports all of the following traffic flows:

- * known unicast traffic from Root to Roots & Leafs
- * known unicast traffic from Leaf to Roots
- * BUM traffic from Root to Roots & Leafs
- * BUM traffic from Leaf to Roots

A particular E-Tree service may need to support all of the above types of flows or only a select subset, depending on the target application. In the case where only multicast and broadcast flows need to be supported, the L2VPN PEs can avoid performing any MAC learning function.

The following subsections will describe the operation of EVPN to support E-Tree service with and without MAC learning.

5.5.1. E-Tree with MAC Learning

The PEs implementing an E-Tree service must perform MAC learning when unicast traffic flows must be supported among Root and Leaf sites. In this case, the PE(s) with Root sites performs MAC learning in the data path over the ESs and advertises reachability in EVPN MAC/IP Advertisement routes. These routes will be imported by all PEs for that BD (i.e., PEs that have Leaf sites as well as PEs that have Root sites). Similarly, the PEs with Leaf sites perform MAC learning in the data path over their ESs and advertise reachability in EVPN MAC/IP Advertisement routes. PEs with Root and/or Leaf sites may use the Ethernet Auto-Discovery per EVI (EAD-EVI) routes for aliasing (in the case of multihomed segments) and EAD-ES routes for mass MAC withdrawal per [RFC7432].

To support multicast/broadcast from Root to Leaf sites, either a P2MP tree rooted at the PE(s) with the Root site(s) (e.g., Root PEs) or Ingress Replication can be used (see Section 16 of [RFC7432]). The multicast tunnels are set up through the exchange of the EVPN Inclusive Multicast route, as defined in [RFC7432].

To support multicast/broadcast from Leaf to Root sites, either Ingress Replication tunnels from each Leaf PE or a P2MP tree rooted at each Leaf PE can be used. The following two paragraphs describe when each of these tunneling schemes can be used and how to signal them.

When there are only a few Root PEs with small amount of multicast/broadcast traffic from Leaf PEs toward Root PEs, then Ingress Replication tunnels from Leaf PEs toward Root PEs should be sufficient. Therefore, if a Root PE needs to support a P2MP tunnel in the transmit direction from itself to Leaf PEs, and, at the same time, it wants to support Ingress Replication tunnels in the receive direction, the Root PE can signal it efficiently by using a new composite tunnel type defined in Section 7.2. This new composite tunnel type is advertised by the Root PE to simultaneously indicate a P2MP tunnel in the transmit direction and an Ingress Replication tunnel in the receive direction for the BUM traffic.

If the number of Root PEs is large, P2MP tunnels (e.g., Multipoint LDP (mLDP) or RSVP-TE) originated at the Leaf PEs may be used; thus, there will be no need to use the modified PMSI Tunnel attribute and the composite tunnel type values defined in Section 7.2.

5.5.2. E-Tree without MAC Learning

The PEs implementing an E-Tree service need not perform MAC learning when the traffic flows between Root and Leaf sites are mainly multicast or broadcast. In this case, the PEs do not exchange EVPN MAC/IP Advertisement routes. Instead, the Inclusive Multicast Ethernet Tag route is used to support BUM traffic. In such scenarios, the small amount of unicast traffic (if any) is sent as part of BUM traffic.

The fields of this route are populated per the procedures defined in [RFC7432], and the multicast tunnel setup criteria are as described in the previous section.

Just as in the previous section, if the number of Root PEs are only a few and, thus, Ingress Replication is desired from Leaf PEs to these Root PEs, then the modified PMSI attribute and the composite tunnel type values defined in Section 7.2 should be used.

5.6. Adaptive Ingress and Egress Filtering of BUM Traffic

It was noted previously that BUM procedure for scenario-2 can be further optimized by performing ingress filtering of BUM traffic from a leaf PE to other leaf-only PEs in case of ingress replication. Furthermore, it was noted that a PE role (for a given BD) as a leaf-only, root-only, or both leaf-and-root can change dynamically as new ACs are added or existing ACs are modified or deleted. Therefore, if such optimization is desired, it must be done dynamically via EVPN signaling extensions and without any operator manual intervention. This section describes the procedures for such an adaptive ingress and egress filtering of BUM traffic when ingress replication is used.

5.6.1. Control Plane Operation: Advertising PE

This section describes control plane procedure for a PE advertising EVPN IMET route used for BUM traffic. The IMET route is specified in [RFC7432] and it carries PMSI Tunnel attribute for identifying tunnel type (i.e., Ingress Replication, PIM-SM P2MP, mLDP P2MP, etc). The IMET route also carries Tunnel Encapsulation extended community as specified in [RFC9012] which identifies encapsulation type over multicast tunnel (i.e., VxLAN, NVGRE, GENEVE, MPLS, etc). In case of non-MPLS overlay encapsulation (e.g., VxLAN, GENEVE) with domain-wide VNI, the advertising PE is also the ingress PE for BUM traffic. However, in case of non-MPLS overlay encapsulation with locally-assigned VNI or MPLS overlay encapsulation with local MPLS label, the advertising PE is the ingress PE for P2MP tunnel type and it is the egress PE for Ingress Replication tunnel type. In addition to PMSI Tunnel Attribute and Tunnel Encapsulation EC, the IMET route is advertised with E-Tree EC in order to support adaptive ingress and egress filtering as highlighted below.

- * For a given BD on a PE, if there are ONLY root ACs but no leaf AC, then no E-Tree EC needs to be advertised and the processing at advertising and receiving PEs are based on [RFC7432]
- * For a given BD on a PE, when the first leaf AC becomes active (multi-homed or single-homed), the PE checks to see if there is any root AC active (multi-homed or single-homed), if there is no root AC active, then the PE (re)-advertises the corresponding IMET route along with E-Tree EC with Root-Indication=0, Leaf-Indication=1, and a valid Leaf VNI.

- * For a given BD on a PE, when the first leaf AC becomes active (multi-homed or single-homed), the PE checks to see if there is any root AC active (multi-homed or single-homed), if there is/are active root AC(s), the PE (re)-advertises the corresponding IMET route along with E-Tree EC with Root-Indication=1, Leaf-Indication=1, and a valid Leaf VNI.
- * If root ACs become active after readvertisement of IMET route with only Leaf-Indication set, then the PE MUST readvertise IMET route with both Root-Indication and Leaf-Indication set along with a valid Leaf VNI
- * If the last root AC becomes inactive and there are only leaf ACs, then the PE readvertises IMET route with only Leaf-Indication set and a valid Leaf VNI in the E-Tree extended community
- * If the last leaf AC become inactive and there are only root ACs, then the PE readvertises IMET route without E-Tree EC

Condition	Root AC	Leaf AC	E-Tree EC Fields
1	No	No	No E-Tree EC
2	Yes	No	No E-Tree EC
3	No	Yes	Root-Flag=0, Leaf-Flag=1 Leaf-VNI= valid or Leaf-VNI= 0xFFFFFFFF
4	Yes	Yes	Root-Flag=1, Leaf-Flag=1 Leaf-VNI= valid or Leaf-VNI= 0xFFFFFFFF

Figure 4: E-Tree EC Setting by an Ingress PE per BD

5.6.2. Control Plane Operation: Receiving PE

This section describes control plane procedure for a PE receiving EVPN IMET route used for BUM traffic.

- * For a given BD, if a PE receives IMET route without an E-Tree EC, then the receiving PE treats the peer PE as a root PE and follows the procedure of [RFC7432] and [RFC8365] by adding the advertising PE to the flood list for that BD for ingress replication.

- * For a given BD, if a PE receives IMET route with E-Tree EC that has only Leaf-Indication set and with a valid VNI, then:
 - If the receiving PE is Root-only or Root-and-Leaf, and it is configured for ingress-replication tunnel for that BD, then it adds the advertising PE to its "All-PEs" flood list and excludes the advertising PE from its "non-Leaf" flood list. If the receiving PE is configured for P2MP tunnel for that BD, it joins that tree. The receiving PEs also configure their leaf ACs (if not already configured) for egress filtering of BUM traffic matching Leaf VNI.
 - If the receiving PE is Leaf-only, and it is configured for ingress-replication tunnel for that BD, it excludes the advertising PE from its flood list. If the receiving PE is configured for P2MP tunnel for that BD, it does not join that tree.
- * For a given BD, if a PE receives IMET route with E-Tree EC that has both Root-Indication and Leaf-Indication set along with a valid Leaf VNI, then the receiving PE (Root-only, Root-and-Leaf, or Leaf-only) add the advertising PE to their "All-PEs" flood list if configured for ingress replication tunnel, or join the multicast tree if configured for P2MP tunnel for that BD. The receiving PEs also configure their leaf ACs (if not already configured) for egress filtering of BUM traffic matching Leaf VNI.

5.6.3. Data Plane Operation

- * For ingress replication:
 - BUM traffic coming from a leaf AC on the local PE uses "non-Leaf" flood list
 - BUM traffic coming from a root AC uses "All-PEs" flood list
- * For P2MP or MP2MP tunnels:
 - Leaf-only PEs, do not join the multicast tunnels from other Leaf PEs
 - Root-only or Root-and-Leaf PEs join the multicast tunnels from Leaf PEs
 - All PEs join the multicast tunnels from Root-only and Root-and-Leaf PEs
- * For all multicast tunnels, imposition PEs:

- BUM traffic coming from a leaf AC on the local PE uses leaf VNI
- BUM traffic coming from a root AC on the local PE uses baseline VNI

* For all multicast tunnels, disposition PEs:

- BUM traffic received with leaf VNI is only sent to root ACs
- BUM traffic received with base VNI is sent to both root and leaf ACs per baseline procedure

Condition	Ingress PE role	Ingress AC role	Flood List
1	Root	Root	Use All-PEs flood list with Base VNI for BUM traffic
2	Leaf	Leaf	Use non-Leaf-PEs flood list with Leaf VNI for BUM traffic
3	Root+ Leaf	Root	Use All-PEs flood list with Base VNI for BUM traffic
4	Root+ Leaf	Leaf	Use non-Leaf-PEs flood list with Leaf VNI for BUM traffic

Figure 5: When to use "non-Leaf PEs" flood list

6. Operation for PBB-EVPN

In PBB-EVPN, the PE advertises a Root or Leaf-Indication along with each Backbone MAC (B-MAC) Advertisement route to indicate whether the associated B-MAC address corresponds to a Root or a Leaf site. Just like the EVPN case, the new E-Tree extended community defined in Section 7 is advertised with each EVPN MAC/IP Advertisement route.

In the case where a multihomed ES has both Root and Leaf sites attached, two B-MAC addresses are advertised: one B-MAC address is per ES (as specified in [RFC7623]) and implicitly denotes Root, and the other B-MAC address is per PE and explicitly denotes Leaf. The former B-MAC address is not advertised with the E-Tree extended community, but the latter B-MAC denoting Leaf is advertised with the new E-Tree extended community where a "Leaf-indication" flag is set. In multihoming scenarios where an ES has both Root and Leaf ACs, it is assumed that while different ACs (VLANs) on the same ES could have

a different Root/Leaf designation (some being Roots and some being Leafs), the same VLAN does have the same Root/Leaf designation on all PEs on the same ES. Furthermore, it is assumed that there is no forwarding among subnets (i.e., the service is L2 and not IRB). An IRB use case is outside the scope of this document.

The ingress PE uses the right B-MAC source address depending on whether the Ethernet frame originated from the Root or Leaf AC on that ES. The mechanism by which the PE identifies whether a given frame originated from a Root or Leaf site on the segment is based on

the Ethernet Tag associated with the frame. Other mechanisms of identification, beyond the Ethernet Tag, are outside the scope of this document.

Furthermore, a PE advertises two special global B-MAC addresses, one for Root and another for Leaf, and tags the Leaf one as such in the MAC Advertisement route. These B-MAC addresses are used as source addresses for traffic originating from single-homed segments. The B-MAC address used for indicating Leaf sites can be the same for both single-homed and multihomed segments.

6.1. Known Unicast Traffic

For known unicast traffic, the PEs perform ingress filtering: on the ingress PE, the Customer/Client MAC (C-MAC) [RFC7623] destination address lookup yields, in addition to the target B-MAC address and forwarding adjacency, a flag that indicates whether the target B-MAC is associated with a Root or a Leaf site. The ingress PE also checks the status of the originating site; if both are Leafs, then the packet is not forwarded.

6.2. BUM Traffic

For BUM traffic, the PEs must perform egress filtering. When a PE receives an EVPN MAC/IP Advertisement route (which will be used as a source B-MAC for BUM traffic), it updates its egress filtering (based on the source B-MAC address) as follows:

- * If the EVPN MAC/IP Advertisement route indicates that the advertised B-MAC is a Leaf, and the local ES is a Leaf as well, then the source B-MAC address is added to its B-MAC list used for egress filtering (i.e., to block traffic from that B-MAC address). Otherwise, the B-MAC filtering list is not updated.
- * If the EVPN MAC/IP Advertisement route indicates that the advertised B-MAC has changed its designation from a Leaf to a Root, and the local ES is a Leaf, then the source B-MAC address is

removed from the B-MAC list corresponding to the local ES used for egress filtering (i.e., to unblock traffic from that B-MAC address).

When the egress PE receives the packet, it examines the B-MAC source address to check whether it should filter or forward the frame. Note that this uses the same filtering logic as the split-horizon filtering described in Section 6.2.1.3 of [RFC7623] and does not require any additional flags in the data plane.

Just as in Section 5.2, the PE places all Leaf ESs of a given bridge domain in a single split-horizon group in order to prevent intra-PE forwarding among Leaf segments. This split-horizon function applies to BUM traffic as well as known unicast traffic.

6.3. E-Tree without MAC Learning

In scenarios where the traffic of interest is only multicast and/or broadcast, the PEs implementing an E-Tree service do not need to do any MAC learning. In such scenarios, the filtering must be performed on egress PEs. For PBB-EVPN, the handling of such traffic is per Section 6.2 without the need for C-MAC learning (in the data plane) in the I-component (C-bridge table) of PBB-EVPN PEs (at both ingress and egress PEs).

7. BGP Encoding

This document defines a new BGP extended community for EVPN.

7.1. E-Tree Extended Community

This extended community is a new transitive extended community [RFC4360] having a Type field value of 0x06 (EVPN) and the Sub-Type 0x05. It is used for Leaf-Indication of known unicast and BUM traffic. It indicates that the frame is originated from a Leaf site.

The E-Tree extended community is encoded as an 8-octet value as follows:

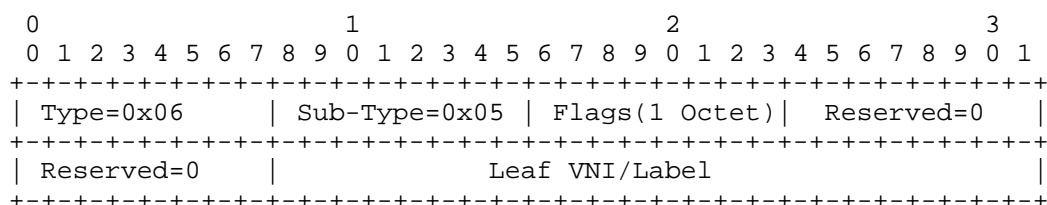


Figure 6: E-Tree Extended Community

The Flags field has the following format:

```

    0 1 2 3 4 5 6 7
    +-----+
    |      MBZ      |R|L|      (MBZ = MUST Be Zero)
    +-----+
```

This document defines the following flags:

- + Leaf-Indication (L)
- + Root-Indication (R)

A value of one for L flag indicates a Leaf AC/site.

A value of one for R flag indicates a Root AC/site.

The rest of the flag bits (MBZ field) are reserved and must be set to zero.

When this extended community is advertised along with the MAC/IP Advertisement route (for known unicast traffic) per Section 5.1, the Leaf-Indication flag MUST be set to one and the Leaf label SHOULD be set to zero. The receiving PE MUST ignore Leaf label and the Root-Indication flag, and only process the Leaf-Indication flag. A value of zero for the Leaf- Indication flag is invalid when sent along with a MAC/IP Advertisement route, and an error should be logged.

When this extended community is advertised along with the EAD-ES route (with an ESI of zero) for BUM traffic to enable egress filtering on disposition PEs per Sections 4.2.1 and 4.2.3, the Leaf label MUST be set to a valid MPLS label (i.e., a non-reserved, assigned MPLS label [RFC3032]) and the Leaf-Indication flag SHOULD be set to zero. The value of the 20-bit MPLS label is encoded in the high-order 20 bits of the Leaf label field. The receiving PE MUST ignore the Leaf-Indication and the Root-Indication flags. A non-valid MPLS label, when sent along with the EAD-ES route, should be ignored and logged as an error.

When this extended community is advertised along with the IMET route, then the procedures covered in "Operation for EVPN" must be followed.

The reserved bits SHOULD be set to zero by the transmitter and MUST be ignored by the receiver.

7.2. PMSI Tunnel Attribute

[RFC6514] defines the PMSI Tunnel attribute, which is an optional transitive attribute with the following format:

=====		=====
	Flags (1 octet)	
=====		=====
	Tunnel Type (1 octet)	
-----		-----
	Ingress Replication MPLS Label (3 octets)	
-----		-----
	Tunnel Identifier (variable)	
-----		-----

Table 1: PMSI Tunnel Attribute

This document defines a new composite tunnel type by introducing a new 'composite tunnel' bit in the Tunnel Type field and adding an MPLS label to the Tunnel Identifier field of the PMSI Tunnel attribute, as detailed below. All other fields remain as defined in [RFC6514]. Composite tunnel type is advertised by the Root PE to simultaneously indicate a non-Ingress-Replication tunnel (e.g., P2MP tunnel) in the transmit direction and an Ingress Replication tunnel in the receive direction for the BUM traffic.

When receiver Ingress Replication labels are needed, the high-order bit of the Tunnel Type field (composite tunnel bit) is set while the remaining low-order seven bits indicate the Tunnel Type as before (for the existing Tunnel Types). When this composite tunnel bit is set, the "tunnel identifier" field begins with a three-octet label, followed by the actual tunnel identifier for the transmit tunnel. PEs that don't understand the new meaning of the high-order bit treat the Tunnel Type as an undefined Tunnel Type and treat the PMSI Tunnel attribute as a malformed attribute [RFC6514]. That is why the composite tunnel bit is allocated in the Tunnel Type field rather than the Flags field. For the PEs that do understand the new meaning of the high-order, if Ingress Replication is desired when sending BUM traffic, the PE will use the label in the Tunnel Identifier field when sending its BUM traffic.

Using the composite tunnel bit for Tunnel Types 0x00 'no tunnel information present' and 0x06 'Ingress Replication' is invalid. A PE that receives a PMSI Tunnel attribute with such information considers it malformed, and it SHOULD treat this Update as though all the routes contained in this Update had been withdrawn per Section 6 of [RFC6514].

8. Security Considerations

Since this document uses the EVPN constructs of [RFC7432] and [RFC7623], the same security considerations in these documents are also applicable here. Furthermore, this document provides an additional security check by allowing sites (or ACs) of an EVPN instance to be designated as a "Root" or "Leaf" by the network operator / service provider and thus prevent any traffic exchange among "Leaf" sites of that VPN through ingress filtering for known unicast traffic and egress filtering for BUM traffic. Since (by default and for the purpose of backward compatibility) an AC that doesn't have a Leaf designation is considered a Root AC, in order to avoid any traffic exchange among Leaf ACs, the operator SHOULD configure the AC with a proper role (Leaf or Root) before activating the AC.

9. IANA Considerations

IANA has allocated sub-type value 5 in the "EVPN Extended Community Sub-Types" registry defined in [RFC7153] as follows:

SUB-TYPE VALUE	NAME	Reference
-----	-----	-----
0x05	E-Tree Extended Community	[RFC8317]

This document creates a one-octet registry called "E-Tree Flags". New registrations will be made through the "RFC Required" procedure defined in [RFC8126]. Initial registrations are as follows:

Bit	Name	Reference
----	-----	-----
0-5	Unassigned	
6	Root-Indication	[This document]
7	Leaf-Indication	[RFC8317]

The Root-Indication flag is sent only along EVPN IMET route and not Eth-AD per ES or MAC/IP routes

9.1. Considerations for PMSI Tunnel Types

The "P-Multicast Service Interface (PMSI) Tunnel Types" registry in the "Border Gateway Protocol (BGP) Parameters" registry has been updated to reflect the use of the most significant bit as the "composite tunnel" bit (see Section 7.2).

For this purpose, this document updates [RFC7385] by changing the previously unassigned values (i.e., 0x08 - 0xFA) as follows:

Value	Meaning	Reference
-----	-----	-----
0x0C-0x7A	Unassigned	
0x7B-0x7E	Experimental	[RFC8317]
0x7F	Reserved	[RFC8317]
0x80-0xFA	Reserved for Composite Tunnel	[RFC8317]
0xFB-0xFE	Experimental	[RFC7385]
0xFF	Reserved	[RFC7385]

The allocation policy for values 0x08-0x7A is per IETF Review [RFC8126]. The range for "Experimental" has been expanded to include the previously assigned range of 0xFB-0xFE and the new range of 0x7B-0x7E. The values in these ranges are not to be assigned. The value 0x7F, which is the mirror image of (0xFF), is reserved in this document.

10. Additional Authors from RFC8317

Samer Salam, Cisco

Email: ssalam@cisco.com

James Uttaro, ATT

Email: uttaro@att.com

Sami Boutros, Ciena

Email: sboutros@ciena.com

Acknowledgements

For [RFC8317], we would like to thank Eric Rosen, Jeffrey Zhang, Wen Lin, Aldrin Issac, Wim Henderickx, Dennis Cai, and Antoni Przygienda for their valuable comments and contributions. The authors would also like to thank Thomas Morin for shepherding this document and providing valuable comments.

For this Document, we would like to thank Neeraj Malhotra, Ramchander Nadipally, Lukas Krattiger, Arie Vayner, Akhil Shashidhar, and Sergey Kolobov for their valuable inputs.

References

Normative References

- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <<https://www.rfc-editor.org/info/rfc1997>>.
- [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>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
- [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>>.
- [RFC7153] Rosen, E. and Y. Rekhter, "IANA Registries for BGP Extended Communities", RFC 7153, DOI 10.17487/RFC7153, March 2014, <<https://www.rfc-editor.org/info/rfc7153>>.
- [RFC7385] Andersson, L. and G. Swallow, "IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points", RFC 7385, DOI 10.17487/RFC7385, October 2014, <<https://www.rfc-editor.org/info/rfc7385>>.
- [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>>.
- [RFC7623] Sajassi, A., Ed., Salam, S., Bitar, N., Isaac, A., and W. Henderickx, "Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)", RFC 7623, DOI 10.17487/RFC7623, September 2015, <<https://www.rfc-editor.org/info/rfc7623>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

- [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>>.
- [RFC8317] Sajassi, A., Ed., Salam, S., Drake, J., Uttaro, J., Boutros, S., and J. Rabadan, "Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) and Provider Backbone Bridging EVPN (PBB-EVPN)", RFC 8317, DOI 10.17487/RFC8317, January 2018, <<https://www.rfc-editor.org/info/rfc8317>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.

Informative References

- [I-D.ietf-bess-bgp-srv6-args]
Talaulikar, K., Raza, S. K., Rabadan, J., and W. Lin, "SRv6 Argument Signaling for BGP Services", Work in Progress, Internet-Draft, draft-ietf-bess-bgp-srv6-args-02, 3 November 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-bgp-srv6-args-02>>.
- [I-D.ietf-bess-mvpn-evpn-aggregation-label]
Zhang, Z. J., Rosen, E. C., Lin, W., Li, Z., and I. Wijnands, "MVPN/EVPN Tunnel Aggregation with Common Labels", Work in Progress, Internet-Draft, draft-ietf-bess-mvpn-evpn-aggregation-label-14, 4 October 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-bess-mvpn-evpn-aggregation-label-14>>.
- [I-D.ietf-nvo3-vxlan-gpe]
Maino, F., Kreeger, L., and U. Elzur, "Generic Protocol Extension for VXLAN (VXLAN-GPE)", Work in Progress, Internet-Draft, draft-ietf-nvo3-vxlan-gpe-13, 4 November 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-nvo3-vxlan-gpe-13>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.

- [RFC7387] Key, R., Ed., Yong, L., Ed., Delord, S., Jounay, F., and L. Jin, "A Framework for Ethernet Tree (E-Tree) Service over a Multiprotocol Label Switching (MPLS) Network", RFC 7387, DOI 10.17487/RFC7387, October 2014, <<https://www.rfc-editor.org/info/rfc7387>>.
- [RFC7796] Jiang, Y., Ed., Yong, L., and M. Paul, "Ethernet-Tree (E-Tree) Support in Virtual Private LAN Service (VPLS)", RFC 7796, DOI 10.17487/RFC7796, March 2016, <<https://www.rfc-editor.org/info/rfc7796>>.
- [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>>.
- [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>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9135] Sajassi, A., Salam, S., Thoria, S., Drake, J., and J. Rabadan, "Integrated Routing and Bridging in Ethernet VPN (EVPN)", RFC 9135, DOI 10.17487/RFC9135, October 2021, <<https://www.rfc-editor.org/info/rfc9135>>.

Authors' Addresses

Ali Sajassi (editor)
Cisco
Email: sajassi@cisco.com

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

John Drake
Independent
Email: je_drake@yahoo.com

Arivudainambi Appachi gounder
Google
Email: aappachi@google.com

Aaron Bamberger
Arista Networks
Email: abamberger@arista.com