Workgroup: BESS Workgroup
Internet-Draft:
draft-rabadan-bess-evpn-inter-domain-opt-b-03
Published: 4 March 2024
Intended Status: Informational
Expires: 5 September 2024
Authors: J. Rabadan, Ed. Nokia
S. Sathappan Nokia
A. Sajassi Cisco
W. Lin Juniper

EVPN Inter-Domain Option-B Solution

Abstract

An EVPN Inter-Domain interconnect solution is required if two or more sites of the same Ethernet Virtual Private Network (EVPN) are attached to different IGP domains or Autonomous Systems (AS), and they need to communicate. The Inter-Domain Option-B connectivity model is one of the most popular solutions for such EVPN connectivity. While multiple documents refer to this type of interconnect solution and specify different aspects of it, there is no document that summarizes the impact of the Inter-Domain Option-B connectivity in the EVPN procedures. This document does not specify new procedures but analyses the EVPN procedures in an Inter-Domain Option-B network and suggests potential solutions for the described issues. Those solutions are based on either other specifications or based on local implementations that do not modify the end-to-end EVPN control plane.

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 5 September 2024.

Table of Contents

1. Introduction

An EVPN Inter-Domain interconnect solution is required if two or more sites of the same Ethernet Virtual Private Network (EVPN) [I-D.ietf-bess-rfc7432bis] are attached to different IGP domains or Autonomous Systems (AS), and they need to communicate. In general, there are different types of EVPN Inter-Domain models that are classified depending on the procedures implemented on the Border Routers interconnecting the domains. The industry typically classifies the models into three groups:

* EVPN Service Interworking Solution: also referred to as the Service Gateway solution, since the Border Routers instantiate Virtual Routing and Forwarding tables (MAC-VRFs and/or IP-VRFs) and perform a lookup (after decapsulating the transport headers) on those tables so that packets are forwarded between domains. [RFC9014], [I-D.sr-bess-evpn-vpws-gateway] and [I-D.ietf-bess-evpn-ipvpn-interworking] specify the Service Gateway solution for EVPN ELAN, VPWS and Layer-3 services, respectively.

* Inter-Domain Option-B Solution: described in [RFC8365] section 10, this solution provides an interconnect solution for EVPN services by using Border Routers that re-write the EVPN BGP next hops and program a swap operation of the VNIs or MPLS labels (depending on whether the encapsulation is NVO-based or MPLS-based). The "Option-B" term refers to the resemblance of this model with the Multi-AS "type B" interconnect for IP-VPN in [RFC4364], only that this document uses the model for the EVPN family. This solution does not require the instantiation of Virtual Routing and Forwarding tables (VRFs) on the Border Routers.

* Inter-Domain Transport Solution: refers to any Inter-Domain solution that provides connectivity at the transport layer, and therefore does not instantiate VRFs or re-write EVPN BGP next hops or programs swap operations of the EVPN service identifiers (such as VNIs or MPLS service labels) on the Border Routers. The Inter-AS Option-C model described in [RFC4364] section 10 subsection "c" (only that the procedures would be used for EVPN routes, as opposed to VPN-IPv4 and VPN-IPv6 routes in [RFC4364]) is an example of Inter-Domain Transport Solution.

The Inter-Domain Option-B connectivity model is one of the most popular solutions for Inter-Domain EVPN connectivity, due to the fact that it provides isolation for each of the interconnected domains (it prevents the need to leak PE loopbacks between domains) while it does not require the instantiation of VRFs on the Border Routers. While multiple documents refer to this type of interconnect

solution and specify different aspects of it, there is no document that summarizes the impact of the Inter-Domain Option-B connectivity in the EVPN procedures. This document does not specify new procedures but analyses the EVPN procedures in an Inter-Domain Option-B network for:

- *Multi-Homing
- *EVPN E-Tree
- *BUM and IP Multicast forwarding using Ingress Replication or Point-to-Multi-Point tunnels
- *Other EVPN services and including Network Virtualization Overlay (NVO) encapsulations or MPLS-based encapsulations

and provide some guidelines for the described issues. Those guidelines are based on either other specifications or based on local implementations that do not modify the end-to-end EVPN control plane.

1.1. Terminology and Conventions

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.

- *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 BD, then the Ethernet segment is defined to be operating in All-Active redundancy mode.
- *BD: Broadcast Domain. An EVI may be comprised of one BD (VLAN-based or VLAN Bundle services) or multiple BDs (VLAN-aware Bundle services). This document makes use of the term "BD" as described in [I-D.ietf-bess-evpn-irb-mcast] section 1.1.4.
- *BR: Border Router, router that provides connectivity between domains, typically an Area Border Router (ABR) or Autonomous System Border Router (ASBR).
- *BUM traffic: Broadcast, Unknown unicast and Multicast traffic.
- *CE: Customer Edge device, e.g., a host, router, or switch.
- *DF and non-DF: Designated Forwarder and non Designated Forwarder. In an Ethernet Segment, the Designated Forwarder PE or Service Gateway forwards unicast and BUM traffic. The non-Designated

Forwarder PE or Service Gateway blocks BUM traffic (if working in All-Active redundancy mode) or unicast and BUM (if working in Single-Active redundancy mode).

*E-PE: Egress 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'.

*EVI: An EVPN instance spanning the Provider Edge (PE) devices participating in that EVPN.

*MAC-VRF: A Virtual Routing and Forwarding table for Media Access Control (MAC) addresses on a PE. In VLAN-based or VLAN Bundle modes [I-D.ietf-bess-rfc7432bis] a BD is equivalent to a MAC-VRF.

*MPLS and non-MPLS NVO tunnels: refer to Multi-Protocol Label Switching (or the absence of it) Network Virtualization Overlay tunnels. Network Virtualization Overlay tunnels use an IP encapsulation for overlay frames, where the source IP address identifies the ingress PE (or ingress Border Router) and the destination IP address the egress PE (or egress Border Router).

*I-PE: Ingress PE.

*IP-VRF: A VPN Routing and Forwarding table for IP routes on an PE. In this document, an IP-VRF is an instantiation of a layer 3 EVPN service in a PE as per [RFC9135][RFC9136].

*IRB: Integrated Routing and Bridging

*IRB Interface: Integrated Bridging and Routing Interface. A virtual interface that connects the Bridge Table and the IP-VRF on an NVE.

*PE: Provider Edge device. In this document a PE can be a Leaf node in a Data Center or a traditional Provider Edge router in an MPLS network.

*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 BD, then the Ethernet segment is defined to be operating in Single-Active redundancy mode.

*PMSI: Provider Multicast Service Interface.

*SBD: Supplementary Broadcast Domain, a special BD that has an IRB interface to an IP-VRF and it is used in the Optimized Inter-Subnet Multicast model, as described in [I-D.ietf-bess-evpn-irb-mcast].

*SR-MPLS SID: Segment Routing MPLS Segment IDentifier.

*SRv6 SID: Segment Routing for IPv6 Segment IDentifier.

*VRF: A generic Virtual Routing and Forwarding table, used in this document to indicate the instantiation of an EVPN service onto a PE. This service can be any supported EVPN service such as layer-2 multipoint services [I-D.ietf-bess-rfc7432bis], EVPN VPWS [RFC8214], EVPN E-Tree [RFC8317], PBB-EVPN [RFC7623], or Layer-3 services as defined in [RFC9135] or [RFC9136].

*VPWS: EVPN Virtual Private Wire Service, as in [RFC8214].

2. EVPN Inter-Domain Option-B General Procedures

The EVPN Inter-Domain Option-B procedures are applied in Border Routers that interconnect domains, and the Ingress and Egress PEs should be configured and operated in the same way they are when communicating with other PEs within their domain. The typical deployments are illustrated in Figure 1 and Figure 2. Figure 1 illustrates an Inter-Domain example where each domain is an IGP instance. The Border routers BR-1 and BR-2 show direct BGP EVPN neighboring between them, and also with the Ingress PE (I-PE) and the Egress PE (E-PE) respectively. However, Route Reflectors may exist in each of the domains. The procedures described in this document remain unchanged irrespective of the presence of Route Reflectors in each domain. Note that in this document VRF is generically used, and may mean either MAC-VRF or IP-VRF, unless otherwise specified.

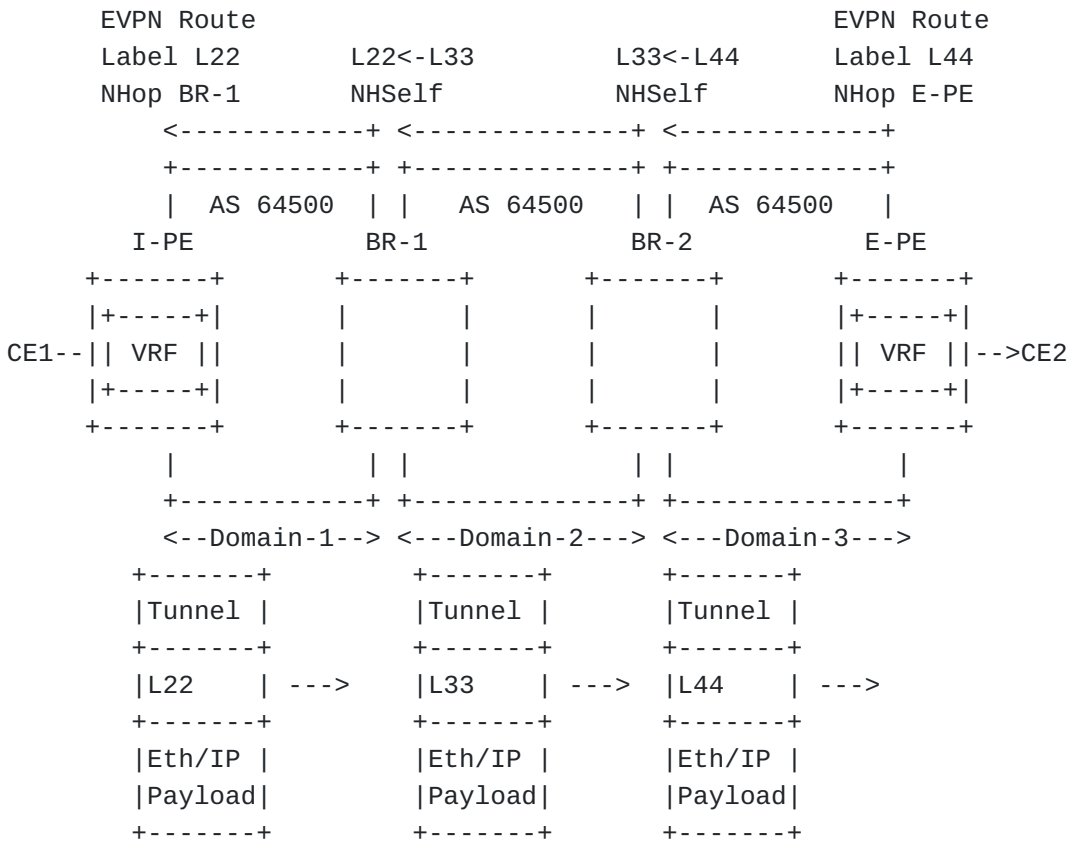

Figure 1: EVPN Inter-Domain Option-B scenario for IGP domains

This document describes also the Inter-Domain Option-B aspects in scenarios such as the one portrayed in Figure 2, where the Border Routers connect different Autonomous Systems. As in the case in Figure 1 the procedures do not change in case the Domains use Route Reflectors.

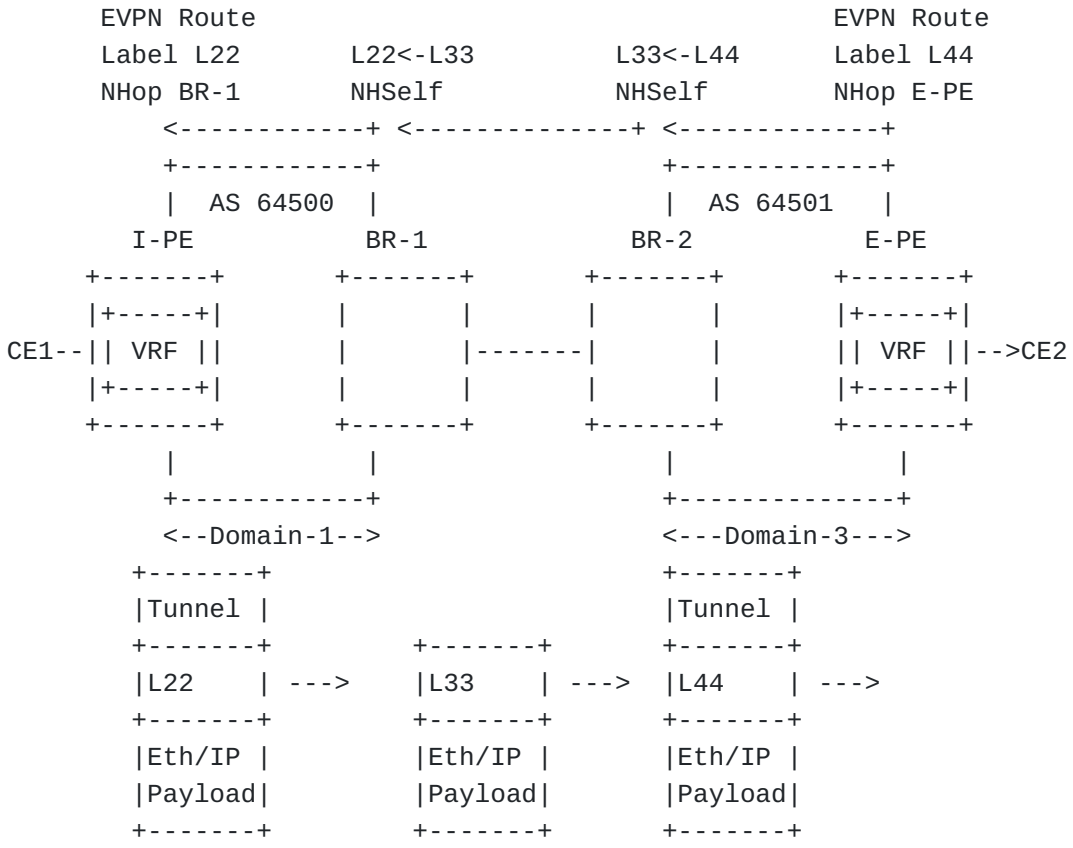

Figure 2: EVPN Inter-Domain Option-B scenario for Multi-AS Backbones

In either Figure 1 or Figure 2, this Inter-Domain Option-B solution involves the redistribution of EVPN routes from domain to domain by the Border Routers. A Border Router learns all the EVPN routes of its own domain, typically via IBGP from the Egress PE or as a client from the domain's Route Reflector, and readvertises those routes to the neighboring Border Router(s), via EBGP or IBGP. When redistributing EVPN routes to the adjacent Border Routers or Route Reflectors within the adjacent domain, the Border Router changes the Next Hop IP address to itself, and the EVPN label of the readvertised BGP MP_REACH_NLRI message to a new generated label. In essence, this means that the Border Router programs a label swap operation in the data path for the EVPN label. For example, packets received on BR-1 with EVPN label L22 are looked up and switched to the interface to the next domain or Border Router, now with EVPN label L33. The EVPN label in this document can be a 20-bit label (that is, an MPLS label or Segment Routing MPLS Segment Identifier) or a 24-bit label (that is, a VNI label for non-MPLS NVO tunnels).

For EVPN routes with 20-bit EVPN labels, in case the Border Router receives the EVPN route via IBGP, the route is resolved to a transport MPLS or SR-MPLS tunnel that provides reachability to the Egress PE or the adjacent Border Router. The imported EVPN route is considered valid and redistributed only in the case the Next Hop is

resolved to such a transport tunnel. In case the Border Router receives the EVPN route via single-hop EBGP, the next hop is resolved to a local interface associated to the next hop, and packets matching the Forwarding Information Base entry for that route are forwarded with a single label in the label stack, i.e. the swapped EVPN label.

In Inter-Domain Option-B scenarios where the transport in the domains is based on NVO tunnels, the EVPN routes advertised from the egress PEs (and redistributed by the Border Routers) use 20-bit labels (in case of MPLS NVO tunnels, e.g., MPLSoGRE) or 24-bit labels (in case if non-MPLS NVO tunnels, e.g., VXLAN). The Border Routers in this case not only swap the label (e.g., VNI) for the NVO packets that they route, but they change the source and destination IP address of the router IP header. When the Border Router forwards packets into an adjacent domain, the outer source IP address of the packets is an IP address of the Border Router. The outer destination IP address is given by the next hop of the EVPN route that created the Forwarding Information Base entry.

The key attributes of the solution are that the Border Routers keep each domain isolated from each other, e.g. BR-2 does not leak E-PE's loopback into other domains, and the Border Routers do not need to have VRFs explicitly configured. The latter aspect also means that the Border Routers need to learn all the EVPN routes within their own domain(s) regardless of the Route Targets, as well as readvertise those to the adjacent domains, possibly selecting a subset of the EVPN routes to be redistributed, via RIB-IN or RIB-OUT policy. The solution does no impose any changes or requirements on the Ingress or Egress PEs, or Route Reflectors. The procedures are solely supported on the Border Routers and should be transparent for the Ingress and Egress PEs.

[RFC8365] section 10.2 is the existing specification for Inter-Domain Option-B in case EVPN uses encapsulations with 20-bit or 24-bit labels, and, in particular for the scenario in Figure 2. This document clarifies that the same procedures and issues apply to the scenario in Figure 1. Although the generic operation of the Border Routers on the received EVPN routes is characterized above, Section 2.1 clarifies the expected behavior on each EVPN route type.

2.1. Border Router procedures on EVPN routes

The Border Router behavior described in Section 2 can be summarized in the following tasks performed on each received EVPN BGP UPDATE:

* The Border Router accepts any EVPN route from the Border Routers and PEs it is connected to (possibly filtering some of the routes via RIB-IN import policies).

*Extracts the EVPN label of each EVPN route, either from the NLRI (Network Layer Reachability Information) or from an attribute included in the BGP UPDATE.

*Programs an EVPN label swap operation in the data path, which switches the extracted EVPN label to a locally generated new EVPN label for the same EVPN route.

*Readvertises the EVPN route (assuming the operation is allowed by policy) with:

a. Next Hop Self, i.e., a new IP address owned by the Border Router itself

b. The locally generated EVPN label for the route

However, there are some subtleties with some EVPN route types that are important to clarify in order to guarantee interoperability across implementations. We differentiate between EVPN Labeled Routes and EVPN Unlabeled Routes.

2.1.1. EVPN Labeled Routes

EVPN Labeled Routes are those that carry EVPN Labels or demultiplexors in the NLRI or an attribute of the BGP UPDATE. If those EVPN Labels are used in the Forwarding Information Base of the Border Router to forward packets between domains, the Label is extracted and added to the Forwarding Information Base associated to a swap operation. If those EVPN Labels are not used to forward packets between domains, but they indicate certain properties of the route, e.g.,: ESI Labels or E-Tree Labels, then the Labels are not extracted, programmed or changed when the route is readvertised. The previous statements MUST be applied to existing and future EVPN route types in Inter-Domain Option-B networks. As an example:

a. Ethernet Auto-Discovery per Ethernet Segment Route (or route type 1 per ES)

   Defined in [I-D.ietf-bess-rfc7432bis], this route signals the multi-homing mode information, as well as the value of the ESI label, encoded in the ESI Label extended community. It is used for fast convergence in case of multi-homed PE failures, via the "Mass Withdraw per Ethernet Segment" procedure. When used with an ESI of zero, the route is used to advertised a Leaf Label in the E-Tree extended community [RFC8317]. The Leaf Label is used by the Ingress PE when forwarding BUM traffic generated from a Leaf Attachment Circuit. Both labels, ESI label and Leaf label, are not used for packet forwarding at the Border Router and therefore the Border Router does not extract them. The Border Router MUST preserve the content of the ESI

label or the E-Tree extended community when readvertising the route to the adjacent domain. Although the next hop self operation is performed on the route by the Border Router, none of the NLRI fields are changed when readvertising the route to the adjacent domain.

b. Ethernet Auto-Discovery per EVPN Instance Route (or route type 1 per EVI)

   Defined in [I-D.ietf-bess-rfc7432bis], this route signals the forwarding information associated to the local EVPN-VPWS Attachment Circuit [RFC8214], and when used with a non-zero ESI, it also performs the Aliasing and Backup procedures for multi-homing in EVPN services. The EVPN label encoded in the NLRI of this route is used when forwarding packets, hence the label must be extracted by the Border Router and programmed in the Forwarding Information Base for a swap operation. Besides the next hop self operation and the new valid label to be encoded in the route, the Border Router does not change any other field of the route. This includes the content of the EVPN Layer-2 Attributes extended community advertised with the route. [RFC8214] section 4 discusses the Inter-domain Option-B solution for EVPN-VPWS.

c. MAC/IP Advertisement Route (or route type 2)

   Defined in [I-D.ietf-bess-rfc7432bis], this route advertises forwarding information for MAC and IP addresses that are used by the Ingress PE to populate the layer-2 Forwarding Information Base, the Address Resolution Protocol or Neighbor Discovery tables [RFC9161] or even the layer-3 Forwarding Information Base [RFC9135]. The route's NLRI contains a mandatory EVPN label, Label1, and an optional Label2. In addition to the next hop self operation, a Border Router that receives a route type 2, with only Label1, needs to extract Label1 from the NLRI, program its value in the Forwarding Information Base, and generate a new valid label that is encoded in Label1 when redistributing the route to the adjacent domain. If the received route type 2 contains a value for both, Label1 and Label2, the Border Router needs to program two separate entries in the Forwarding Information Base (for the value in Label1 and the value in Label2) and generate two valid Label1 and Label2 values. The rest of the information in the route, including EVPN extended communities and Default Gateway extended community, is preserved by the Border Router when readvertising. This method at the Border Router is applied irrespective of the Egress PE using an EVPN label per VRF, EVPN label per Ethernet Segment or EVPN label per MAC address. However, using a label per VRF on the Egress PEs has the least

impact on the Border Routers Forwarding Information Base scale, compared to label per MAC or label per Ethernet Segment.

d. Inclusive Multicast Ethernet Tag Route (or route type 3)

   Also defined in [I-D.ietf-bess-rfc7432bis], this route is used for the auto-discovery of the remote PEs attached to the same Broadcast domain, as well as the creation of the flooding tree used to forward BUM traffic by the PEs attached to the same Broadcast Domain. The route type 3 does not contain any EVPN label in its NLRI. The Provider Tunnel (P-Tunnel) identification is carried in the PMSI Tunnel Attribute. When used for Ingress Replication or Assisted Replication tunnel types, the PMSI Tunnel Attribute contains an EVPN Label (downstream allocated) that is extracted by the Border Router and programmed in the Forwarding Information Base in the same way as for the EVPN labels in the routes above. The Border Router generates a valid new label that is encoded in the PMSI Tunnel Attribute of the route readvertised to the adjacent domain. In addition to the next hop self and label swap operation, the Border Router preserves all the fields in the NLRI (including the Originating Router's IP Address) and the attributes of the routes, including the Tunnel Identifier of the PMSI Tunnel Attribute and the Layer 2 Attributes extended community. When the route type 3 uses a P-Tunnel different than Ingress Replication, the Border Router should carry out the segmentation procedures specified in [I-D.ietf-bess-evpn-bum-procedure-updates].

e. IP Prefix Route (or route type 5)

   Specified in [RFC9136], this route allows the Egress PEs to advertise the IPv4 or IPv6 prefixes that they have learned locally in their IP-VRF. The route's NLRI contains an EVPN label that the Option-B Border Router needs to extract and program in the Forwarding Information Base, along with a label swap operation. Besides the next hop self and generating a new valid EVPN label for the IP Prefix route readvertised to the adjacent domain, the Border Router does not change any of the fields in the NLRI and preserves all the attributes along with the route, including EVPN extended communities.

f. Per-Region I-PMSI A-D Route (or route type 9)

   Used for P-Tunnel Segmentation on Border Routers, its definition and procedures are described in [I-D.ietf-bess-evpn-bum-procedure-updates].

g. S-PMSI A-D Route (or route type 10)

Also defined in [I-D.ietf-bess-evpn-bum-procedure-updates], the Border Router should follow the same procedures as for the Inclusive Multicast Ethernet Tag Route above.

2.1.2. EVPN Unlabeled Routes

Examples or EVPN Unlabeled Routes are:

*Ethernet Segment Route (or route type 4)

*Selective Multicast Ethernet Tag Route (or route type 6)

*Multicast Membership Report Synch Route (or route type 7)

*Multicast Leave Synch Route (or route type 8)

*Leaf Auto-Discovery Route (or route type 11)

The Border Router receiving these routes simply redistributes the routes to the adjacent domain with a next hop of itself, and preserving all the attributes that the routes contain.

3. EVPN Inter-Domain Option-B and Multi-Homing

This section summarizes the issues of the Inter-Domain Option-B associated to EVPN Multi-Homing. Figure 3 illustrates the use of multi-homing in an Inter-Domain Option-B example.

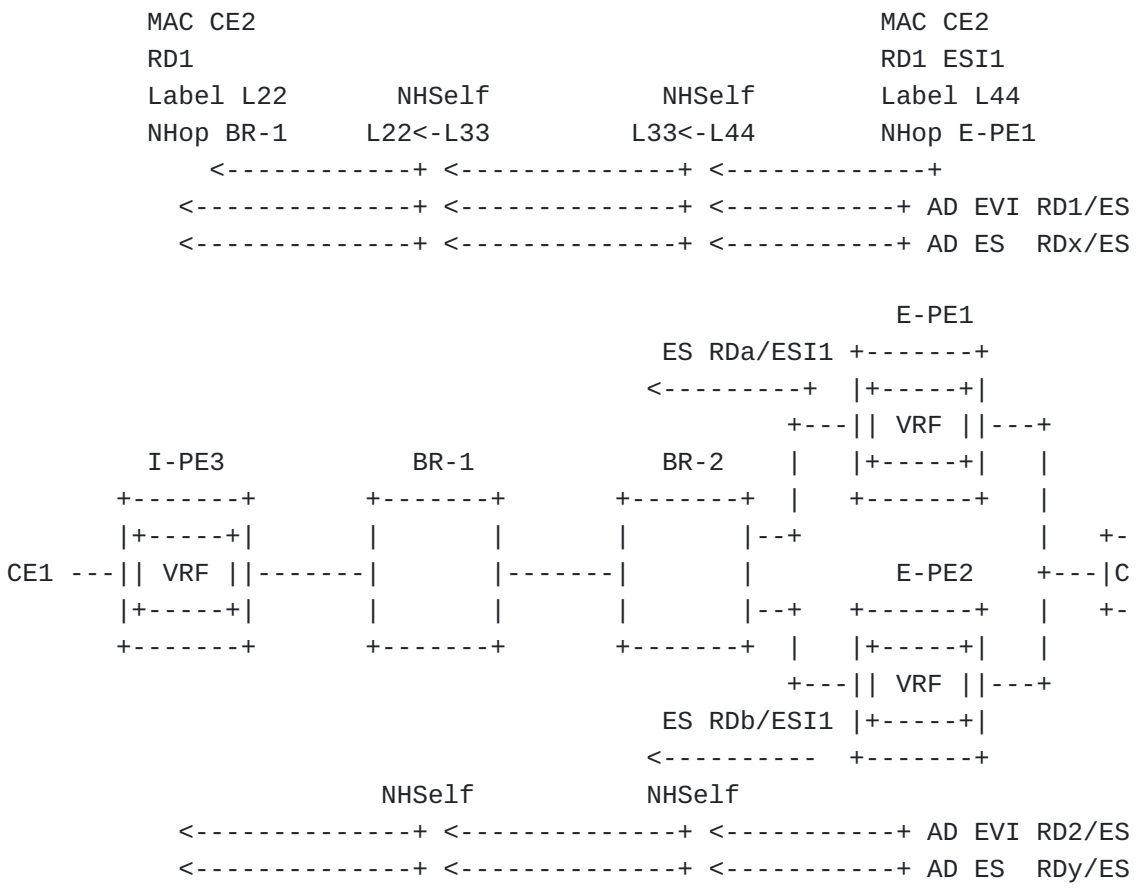

Figure 3: EVPN Inter-Domain Option-B and multi-homing

The Border Router rewriting the EVPN multi-homing routes next hop has an impact on the EVPN multi-homing procedures that follow:

*Mass Withdrawal

*Aliasing and Backup Path procedures

*Designated Forwarder Election and AC-Influenced Capability

*Split Horizon Filtering

3.1. Mass Withdraw

The limitations of the mass withdraw procedures when the multi-homed egress PEs and the ingress PEs are in different domains are explained in [RFC8365] section 10.2.2.

As a refresher, suppose the example of Figure 3 in which CE2 is multi-homed to egress PE1 and PE2 (on Ethernet Segment ES1 with identifier ESI1), and the ingress PE3 sits in a different domain. As illustrated, only E-PE1 advertises the MAC/IP route for MAC CE2, whereas both E-PE1 and E-PE2 advertise the A-D per ES and A-D per EVI routes for ESI1. The fact that the Border Routers rewrite the

next hops of all the routes, prevents I-PE3 from being able to correlate the MAC/IP Advertisement route with the A-D per ES route advertised from the same E-PE, since the only mechanism in [I-D.ietf-bess-rfc7432bis] to correlate A-D per ES and MAC/IP Advertisement routes advertised from the same E-PE is the route next hop. As an example, if the link from CE2 to E-PE1 fails, E-PE1 sends a MP_UNREACH_NLRI message for the A-D per ES route and A-D per EVI route for ESI1. The messages get to I-PE3 and are processed, however, I-PE3 is unable to correlate the withdrawn A-D per ES route with the MAC/IP Advertisement route for CE2 and therefore it does not perform any mass withdraw of the MACs associated to ESI1, as long as at least one A-D per ES route for ESI1 exists. Note that the route distinguisher of the MAC/IP Advertisement route and A-D per ES route advertised from E-PE1 are different, hence the routes cannot be associated.

As also explained in [RFC8365] section 10.2.2, a "mass withdraw per EVI" is possible though, due to the fact that the A-D per EVI routes and MAC/IP Advertisement routes advertised from the same PE and ES can be correlated based on the route distinguisher. In Figure 3, if the link between CE2 and E-PE1 fails, I-PE3 receives the A-D per EVI route withdrawal from E-PE1 and can withdraw all the MACs related to the MAC/IP Advertisement routes that match the route distinguisher of the A-D per EVI route, i.e., RD1 in the example, hence MAC CE2 is flushed on I-PE3. Although the issue is explained for MAC address mass withdrawal, the same issue exists with IP Prefixes, as in [I-D.ietf-bess-evpn-ip-aliasing].

This document assumes that "mass withdraw per EVI" is the default behavior that all PEs and Border Routers MUST support. When "mass withdraw per EVI" is used, unique RDs MUST be used on all the PEs attached to the same EVI.

The following subsections also suggest some potential solutions to overcome the mass withdraw (per ES) limitation imposed by the Border Routers in the Inter-Domain Option-B model. All of them are based on finding a way to correlate the withdrawn A-D per ES route with the routes type 2 and/or 5 advertised by the same egress PE, so that the corresponding MACs or IP Prefixes can be removed.

3.1.1. The Originating PE Attribute Solution

A way to solve the mass withdraw limitation imposed by the Border Routers (for MACs and IP Prefixes) is documented in [I-D.heitz-bess-evpn-option-b], which defines a transitive attribute called Originating PE (OPE) that removes the ambiguity to find the identity of the originator of the routes. When the egress PE advertises the OPE attribute along with the A-D per ES routes and MAC/IP Advertisement or IP Prefix routes, the ingress PE is able to

correlate the routes that are originally advertised from the same egress PE based on the same OPE value received on AD per ES and MAC/IP Advertisement (or IP Prefix) routes. The use of OPE provides a solution to support mass withdrawal per ES in Inter-Domain Option-B networks.

3.1.2. The RD Administrator Subfield Solution

An alternative solution is also hinted by [I-D.heitz-bess-evpn-option-b] section 9.2, where the routes type 2 and 5 can be correlated with the A-D per ES routes from the same PE based on the Administrator subfield of the route distinguishers (RDs). That is, in Figure 3, suppose E-PE1 advertises the A-D per ES route with route distinguisher RDx = <RD1:0> and the MAC/IP Advertisement route with <RD1:1>, with "RD1" being the Administrator subfield of the route distinguisher. E-PE2 allocates "RD2" as Administrator subfield for A-D per ES and MAC/IP Advertisement routes. Now, in case of a withdraw of the A-D per ES route from E-PE1, I-PE3 can perform a mass withdraw operation based on the assumption that all the MACs from the MAC/IP Advertisement routes with RD1 as Administrator subfield are advertised from the same E-PE1 that failed and withdrew the A-D per ES route. The same solution is valid for the mass withdraw of IP Prefix routes.

3.1.3. The EVPN Instance RD Solution

This document suggests a third solution based on the E-PEs using the same route distinguisher on A-D per ES routes and routes type 2 or 5. The A-D per ES routes are normally advertised per <ES, EVI-set>, where an EVI-set is a group of EVPN Instances, each one represented by a different route target in the route. Because of this, the A-D per ES route cannot use the route distinguisher of an existing VRF in the PE, but a unique route distinguisher not assigned to any EVPN Instance (instantiated in a VRF). However, suppose each EVI-set is composed of a single EVI, hence the A-D per ES routes are advertised per <ES, EVI> and therefore there is a separate A-D per ES route per EVPN Instance (or VRF). If that is the case, now the A-D per ES routes can use the route distinguisher assigned to the EVPN Instance (or VRF), which is the same one used by the routes type 2 or 5 for the EVI. Since A-D per ES routes are - with this solution - advertised per <ES, EVI>, this is really a "mass withdraw per EVI" solution, similar to the one described in Section 3.1 in terms of efficiency. However, the advantage of this solution is that the A-D per ES routes are REQUIRED, while A-D per EVI routes are OPTIONAL [I-D.ietf-bess-rfc7432bis] and may not be used in the EVI.

3.2. Aliasing and Backup Path Procedures

The Aliasing and Backup Path procedures work in an Inter-Domain Option-B solution as per [RFC8365], section 10.2. That is, since EVPN MAC/IP Advertisement routes and A-D per EVI routes are both advertised on a per Broadcast Domain basis and they use the same route distinguisher and route target, the receiving ingress PE can associate them together to determine the BGP paths available for the MAC (multiple aliasing paths in case of all-active mode, or one active and one backup in case of single-active mode). Different paths can still be created without ambiguity even if they all go through the same Border Router.

Although the Aliasing and Backup Path procedures per se are not affected, note that the ingress PE installs the MAC from an EVPN MAC/IP Advertisement route (with non-reserved ESI), only if the associated set of Ethernet A-D per ES routes are received from the same egress PE ([I-D.ietf-bess-rfc7432bis], section 9.2.2). Due to the same issues described in Section 3.1, the ingress PE cannot determine if the received MAC/IP Advertisement route and the received set of Ethernet A-D per ES routes are coming from the same egress PE. This document suggests two approaches to solve this resolution issue:

1. Use a "loose" resolution for the MAC/IP Advertisement route - that is, the ingress PE considers the MAC/IP Advertisement route (with a non-reserved ESI) resolved if (and only if) at least one Ethernet A-D per ES route has been received with the same ESI and same next hop as the MAC/IP Advertisement route (it is assumed that its route target set contains the route target of the MAC/IP Advertisement route).

2. Use any of the approaches in Section 3.1 to correlate MAC/IP Advertisement routes and A-D per ES routes, and then resolve the MAC/IP Advertisement route as in ([I-D.ietf-bess-rfc7432bis].

3.3. Designated Forwarder Election and AC-Influenced Capability

On an all-active Ethernet Segment, the Designated Forwarder is the PE router responsible for sending Broadcast, Unknown Unicast, and Multicast (BUM) traffic to a multi-homed Customer Edge (CE) device, in the <ES, Ethernet Tag> for which the PE is elected. If the Ethernet Segment works in single-active mode or port-active mode, the Designated Forwarder is the PE router that sends all traffic to a multi-homed CE [RFC8584]. When a CE is multi-homed to two or more PEs sitting in different domains, the Designated Forwarder candidate list is still created normally. The Designated Forwarder Election is unaffected by the Border Routers next hop self operation on the ES

routes. This is due to the fact that the candidate list is created out of the Originating Router's IP Address of the ES routes (which is not changed by the Border Routers) as opposed to the ES route next hops [RFC8584]. However, the Attachment Circuit Influenced Designated Forwarder (AC-Influenced DF Election) capability [RFC8584] is affected by the next hop self operation of the Border Routers.

If the AC-Influenced DF Election capability is enabled on all the PEs attached to the Ethernet Segment, the Designated Forwarder candidate list needs to be pruned based on the presence of the A-D per ES and A-D per EVI routes for a given candidate. That is, even if E-PE1's ES route is received Figure 3, E-PE2 cannot add E-PE1 to the Designated Forwarder candidate list for <ES1, BD1> until the valid A-D per ES and A-D per EVI routes (for ES1 and BD1) are received and identified as originated from E-PE1. However, because BR-2 changes the next hop of the A-D routes, E-PE2 cannot rely on the next hop to identify the routes as coming from E-PE1. This issue is similar to the one discussed in Section 3.1 for mass withdraw, only that the PE now needs to correlate the ES route and A-D per ES/EVI routes coming from the same PE of origin.

This document assumes that, in case the PEs attached to the same Ethernet Segment are located in different domains, the operator may choose one of the following alternatives:

* Disable the AC-Influenced Designated Forwarder capability in the PEs attached to the Ethernet Segment, or

* Enable the AC-Influenced Designated Forwarder capability in all the PEs attached to the Ethernet Segment, and correlate the received A-D per ES/EVI routes with their corresponding Originating Router's IP Address based on any of the three procedures of Section 3.1.

3.4. Split Horizon Filtering

The Split Horizon Filtering is a fundamental part of the EVPN multi-homing procedures to avoid BUM looped frames to go back to the multi-homed CE. As described in [I-D.ietf-bess-evpn-mh-split-horizon] there are two Split Horizon Filtering Types: ESI label based and Local Bias. Which one is applied depends on the transport tunnel being used by the EVPN BUM packets, and some transport tunnels may support both mechanisms. If two or more PEs of the same Ethernet Segment are sitting in

different domains, the procedures in the Border Router may have an impact on the Split Horizon Filtering mechanisms. In particular:

1. If the multi-homed PEs use an ESI label based Split Horizon Filtering Type:

    a. Regardless of the PEs using upstream or downstream allocated ESI labels (for P2MP/MP2MP or Ingress Replication, respectively), the PEs in the Ethernet Segment need to correlate the identity of the PE advertising the ESI label with the Inclusive Multicast Ethernet Tag routes advertised by the same PE. This brings us back to the same issue of identifying the origin of the A-D per ES route described in Section 3.1, only that this time the receiving PE needs to correlate A-D per ES routes with routes type 3, as opposed to types 2 or 5. In this case, any of the solutions in Section 3.1 could be used.

    b. The use of ESI labels allocated from a Domain-wide Common Block (DCB) and the same label used by all the PEs attached to the same Ethernet Segment may simplify the procedures. If that is the case, the ingress PE can program the received ESI label without the need to correlate the received A-D per ES routes with the Inclusive Multicast Ethernet Tag routes.

    c. In addition, the Border Routers need to preserve the ESI label when they route packets between domains.

2. If the multi-homed PEs use Local Bias as the Split Horizon Filtering Type:

    a. The Border Router cannot change the outer source IP address of the IP tunnel, so that the egress PE can still identify the source PE. Note this may not be possible in many implementations.

The above considerations may influence Inter-Domain Option-B designs, so the capabilities of the Border Routers and PEs have to be analyzed before the operator deploys CEs that are multi-homed to PEs located in different domains.

4. Inter-Domain Option-B and Load Balancing Procedures

This section will cover the impact of Inter-Domain Option-B Border Router procedures in load balancing related mechanisms such as Flow Label or Control Word for MPLS tunnels (see [I-D.ietf-bess-rfc7432bis] section 18), or the source UDP port for NVO tunnels that is used for provide entropy when load balancing

traffic on the core routers. VXLAN [RFC7348] is an example of NVO tunnel type that uses the source UDP port to provide entropy.

4.1. Flow Label

The use of Flow Label and its signaling is described in [I-D.ietf-bess-rfc7432bis] section 18.1. The ingress PE pushes the Flow Label only on EVPN-encapsulated known unicast packets forwarded to egress PEs that previously advertised their Flow Label support on Inclusive Multicast Ethernet Tag routes with the F-bit set. When programming the data path for a given MAC, the ingress PE needs therefore to program the use of Flow Label if the MAC/IP Advertisement route came from the same PE that advertised an Inclusive Multicast Ethernet Tag route with F-bit set. The ingress PE correlates both, MAC/IP Advertisement route and Inclusive Multicast Ethernet Tag route based on the matching route distinguisher of the two.

The Flow Label MUST be preserved by the Border Routers receiving EVPN-encapsulated packets containing a Flow Label, so that the EVPN packets for the same flow are forwarded following the same path within each domain.

4.2. Control Word

The signaling of the Control Word in the Inclusive Multicast Ethernet Tag routes (C-bit) is described in [I-D.ietf-bess-rfc7432bis] section 7.11. As in the case described in Section 4.2, when a Border Router rewrites the next hops of the MAC/IP Advertisement and Inclusive Multicast Ethernet Tag routes, the ingress PE needs to identify the egress PE based on the matching route distinguisher of the two routes. Also, if included in the received EVPN-encapsulated packets, the Control Word MUST be preserved by the Border Routers so that no packet reordering happens for flows forwarded into an adjacent domain.

4.3. Source UDP port

If ingress and egress PEs use NVO tunnels [RFC8365], i.e., IP tunnels, the ingress PE typically encodes a per-flow hash value into the the outer tunnel source UDP port of the EVPN-encapsulated packets. Examples of tunnel types that use the outer source UDP port as an entropy field are VXLAN, GENEVE, or MPLSoUDP. The Border Routers between the ingress and egress PEs MUST preserve the value of the source UDP port so that EVPN-encapsulated packets for the same flow are forwarded following the same path within each domain.

5. Inter-Domain Option-B and Layer-2 MTU

In the same way the support for Flow Label or Control Word is signaled, the egress PE's supported layer-2 MTU (Maximum Transfer Unit) is indicated in the Layer-2 MTU field of the EVPN Layer-2 Attributes extended community advertised along with the Inclusive Multicast Ethernet Tag route ([I-D.ietf-bess-rfc7432bis], section 7.11.1). The Border Router(s) between ingress an egress PEs do not modify any of the advertised attributes, and therefore the layer-2 MTU value is propagated end to end up to the ingress PE. In general, the layer-2 MTU configured in all PEs attached to the same EVPN service SHOULD match, irrespective of the domain where they reside. In case MTUs are different in the different domains, [I-D.ietf-bess-rfc7432bis] allows the signaling a layer-2 MTU of zero from the egress PE, which is not checked at the ingress PE and ensures the EVPN destination is properly programmed at this ingress PE.

6. E-Tree Considerations

[RFC8317], or Ethernet-Tree in EVPN networks, describes two areas that are impacted by the presence of an Inter-Domain Option-B Border Router between ingress and egress PEs: the use of composite tunnels for BUM traffic and the egress PE filtering of BUM traffic originated from a Leaf Attachment Circuit.

6.1. E-Tree Composite Tunnels

A composite tunnel is tunnel type used by the Root PE to simultaneously indicate a P2MP tunnel in the transmit direction and an Ingress Replication tunnel in the receive direction for BUM traffic. For this reason, an Inclusive Multicast Ethernet Tag route for a composite tunnel comprises both, a downstream allocated EVPN label for Ingress replication, and a P2MP tunnel identifier. The EVPN label is extracted by the Border Router and programmed in the Forwarding Information Base, as described in Section 2.1.1 bullet "d". Since the Ingress Replication procedures are followed, the Border Router generates a valid new label that is encoded in the (composite type) PMSI Tunnel Attribute of the route readvertised to the adjacent domain. Also, as described in Section 2.1.1, the segmentation procedures in [I-D.ietf-bess-evpn-bum-procedure-updates] are followed for the encoded P2MP tunnel in the same PMSI Tunnel Attribute.

6.2. Egress Filtering of BUM Traffic Originated from a Leaf Attachment Circuit

E-Tree in EVPN networks requires the filtering of traffic originated from a Leaf Attachment Circuit. While the ingress PE can determine

if known unicast leaf traffic can be forwarded, based on whether the destination MAC address belongs to a leaf Attachment Circuit, filtering of the BUM traffic must be done at the egress PE. For such filtering, the egress PE advertises a Leaf Label along with an Ethernet A-D per ES route (with ESI of zero), and the egress PE relies on the ingress PE to push that Leaf Label when sending Leaf BUM traffic to it [RFC8317]. If ingress and egress PEs are located in different domains of an Inter-Domain Option-B network, the ingress PE cannot correlate the received Inclusive Multicast Ethernet Tag route and A-D per ES route (comprising the Leaf Label) from the same egress PE. Due to this issue when identifying the egress PE's Leaf Label, the ingress PE cannot push the Leaf Label below the EVPN multicast label for a given egress PE. The issue is illustrated in Figure 4.

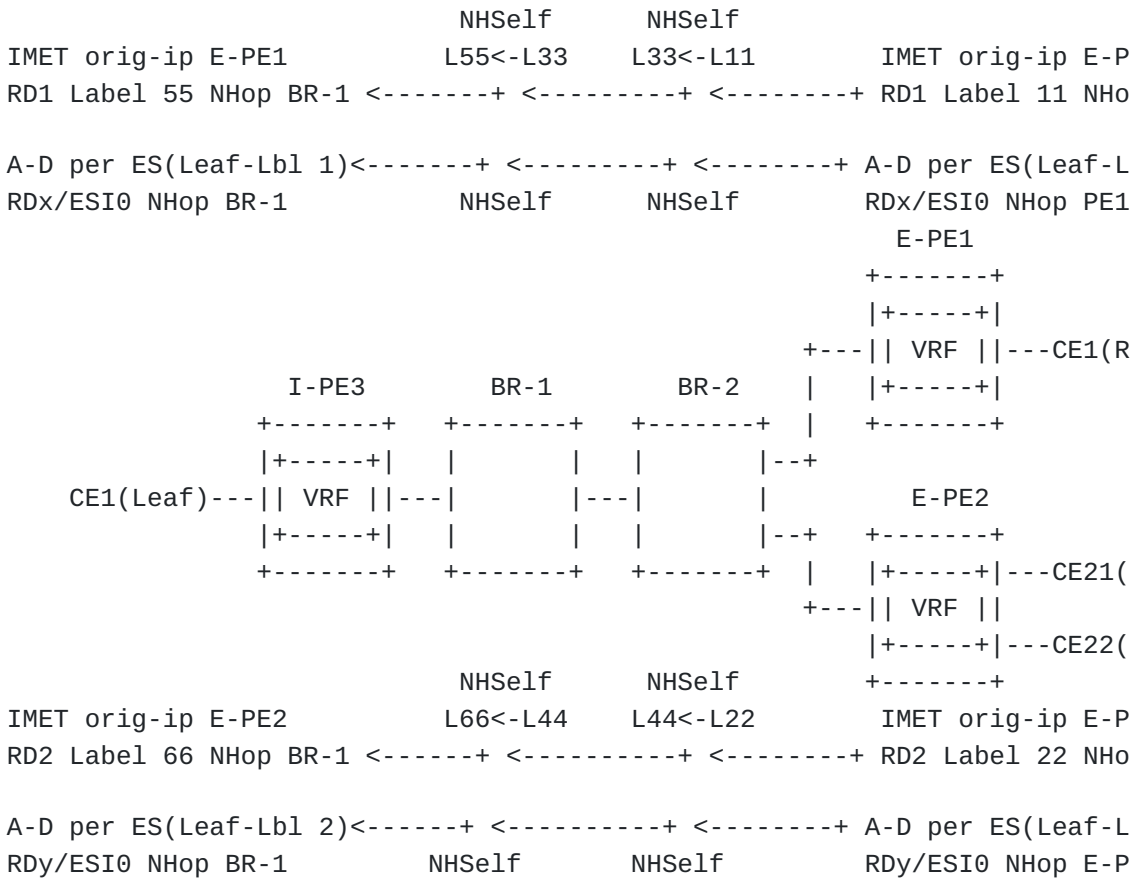

Figure 4: EVPN Inter-Domain Option-B and Leaf BUM filtering

Suppose the egress PEs and ingress PE are in a different domain Figure 4, and that I-PE3 needs to forward EVPN-encapsulated BUM traffic from Leaf CE1, using Ingress Replication. I-PE3 receives Inclusive Multicast Ethernet Tag routes and A-D per ES routes from the two egress PEs, however, I-PE3 is unable to identify what Leaf Label needs to push when sending EVPN-encapsulated BUM traffic to E-

PE1 or E-PE2. This is due to the fact that the A-D per ES routes cannot longer be associated with their corresponding Inclusive Multicast routes based on the next hop, since the four routes in the example are received from the same next hop. This section suggests different solutions, as follows.

6.2.1. Identication of the PE of Origin

A way to solve the issue with E-Tree and the egress filtering of Leaf BUM traffic is to identify and correlate the Inclusive Multicast Ethernet Tag routes and A-D per ES routes (with ESI of zero) originated from the same egress PE. In order to do that, any of the three techniques in Section 3.1 are valid, only that the identification is now done so that Inclusive Multicast Ethernet Tag routes and A-D per ES routes can be correlated, instead of MAC/IP Advertisement routes and A-D per ES routes.

6.2.2. Domain-wide Common Block Leaf Labels

The use of Leaf Labels allocated from a Domain-wide Common Block (DCB) and the same Leaf label value used by all the PEs attached to the E-Tree EVPN service simplify the procedures. If that is the case, all the egress PEs advertise the same Leaf label in their A-D per ES routes for ESI of zero, and that Label value matches the local Leaf label on the ingress PE. The ingress PE can then program the allocated Leaf label for all the destination egress PEs, without correlating the received Inclusive Multicast and A-D per ES routes. This assumes all the PEs in the Broadcast Domain allocate the same Leaf label. If the ingress PE detects any inconsistency in the signaled Leaf label, that is, if at least one PE of the Broadcast Domain advertises a different label than the local Leaf label, then the ingress PE SHOULD NOT program the Leaf label when sending traffic to the egress PEs.

6.2.3. Source MAC-based Egress Filtering

Another potential solution is the use of source MAC-based egress filtering, as opposed to Leaf label-based egress filtering for EVPN-encapsulated BUM traffic. If the ingress PE receives two or more A-D per ES routes (with ESI of zero) with the same next hop, then it does not program any of the received Leaf labels and forwards EVPN-encapsulated BUM packets with the EVPN label and without any Leaf label. If we assume that the ingress PE has previously advertised the local Leaf MAC addresses, when the BUM packets get to the egress PE, a source MAC lookup in the MAC-VRF will determine if the BUM packet is coming from a Leaf or a Root Attachment Circuit.

Taking the example of Figure 4, I-PE3 advertises CE1's MAC as a Leaf MAC in a route type 2, and hence CE1's MAC is programmed in E-PE1

and E-PE2 as Leaf. Since I-PE3 receives two A-D per ES routes (with ESI of zero) from the same next hop, I-PE3 determines that it cannot program the received Leaf labels, and therefore I-PE3 forwards BUM packets from CE1 to E-PE1 and E-PE2 with their corresponding Inclusive Multicast labels and without any Leaf label. When the packets get to the egress PEs, E-PE1 and E-PE2 perform a source MAC lookup in the MAC-VRF. Since CE1's MAC appear as a Leaf MAC, E-PE1 and E-PE2 can filter appropriately. That is, e.g., E-PE2 forwards to CE22 (root) only and not to CE21 (leaf).

7. Inter-Domain Option-B and PBB-EVPN

Provider Backbone Bridging EVPN [RFC7623] is also supported in Inter-Domain Option-B. The following considerations apply:

* PBB-EVPN does not have any of the issues described in Section 3. This is due to the fact that PBB-EVPN multi-homing procedures do not rely on Ethernet A-D per ES or per EVI routes at all.

* PBB-EVPN does not have any of the issues described in Section 6 either, for the same reason. For E-Tree egress filtering of the EVPN-encapsulated BUM packets (so that they are only forwarded to local Root Attachment Circuits and not Leaf Attachment Circuits), PBB-EVPN relies on the source B-MAC identification at the egress PE. The procedures are not impacted by the presence of a Border Router between ingress and egress PEs.

* Also, this document assumes that the [I-D.ietf-bess-rfc7432bis] procedures to signal Flow Label, Control Word or Layer-2 MTU, do not apply to PBB-EVPN networks, hence there are no issues derived from those components.

8. Security Considerations

This document is intended to be published as Informational and hence does not impose and procedures that introduce any new security risks. The described solutions are based on existing specifications and therefore this document inherits the security considerations described in each of the normative reference documents.

9. IANA Considerations

No IANA actions.

10. Contributors

11. Acknowledgments

The authors would like to thank Jeffrey Zhang for his review and comments.

12. References

12.1. 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>.

[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>.

[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>.

[RFC9136] Rabadan, J., Ed., Henderickx, W., Drake, J., Lin, W., and A. Sajassi, "IP Prefix Advertisement in Ethernet VPN (EVPN)", RFC 9136, DOI 10.17487/RFC9136, October 2021, <https://www.rfc-editor.org/info/rfc9136>.

[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>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <https://www.rfc-editor.org/info/rfc4364>.

[I-D.ietf-bess-rfc7432bis] Sajassi, A., Burdet, L. A., Drake, J., and J. Rabadan, "BGP MPLS-Based Ethernet VPN", Work in Progress, Internet-Draft, draft-ietf-bess-rfc7432bis-08, 13 February 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-rfc7432bis-08>.

[RFC9014] Rabadan, J., Ed., Sathappan, S., Henderickx, W., Sajassi, A., and J. Drake, "Interconnect Solution for Ethernet VPN (EVPN) Overlay Networks", RFC 9014, DOI 10.17487/RFC9014, May 2021, <https://www.rfc-editor.org/info/rfc9014>.

[RFC8214] Boutros, S., Sajassi, A., Salam, S., Drake, J., and J. Rabadan, "Virtual Private Wire Service Support in

Ethernet VPN", RFC 8214, DOI 10.17487/RFC8214, August 2017, <https://www.rfc-editor.org/info/rfc8214>.

[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>.

[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>.

[RFC8584] Rabadan, J., Ed., Mohanty, S., Ed., Sajassi, A., Drake, J., Nagaraj, K., and S. Sathappan, "Framework for Ethernet VPN Designated Forwarder Election Extensibility", RFC 8584, DOI 10.17487/RFC8584, April 2019, <https://www.rfc-editor.org/info/rfc8584>.

[I-D.ietf-bess-evpn-irb-mcast] Lin, W., Zhang, Z. J., Drake, J., Rosen, E. C., Rabadan, J., and A. Sajassi, "EVPN Optimized Inter-Subnet Multicast (OISM) Forwarding", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-irb-mcast-10, 27 February 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-irb-mcast-10>.

[I-D.ietf-bess-evpn-ipvpn-interworking] Rabadan, J., Sajassi, A., Rosen, E. C., Drake, J., Lin, W., Uttaro, J., and A. Simpson, "EVPN Interworking with IPVPN", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-ipvpn-interworking-09, 9 October 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-ipvpn-interworking-09>.

12.2. Informative References

[RFC9161] Rabadan, J., Ed., Sathappan, S., Nagaraj, K., Hankins, G., and T. King, "Operational Aspects of Proxy ARP/ND in Ethernet Virtual Private Networks", RFC 9161, DOI 10.17487/RFC9161, January 2022, <https://www.rfc-editor.org/info/rfc9161>.

[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>.

[I-D.ietf-bess-evpn-bum-procedure-updates]
Zhang, Z. J., Lin, W., Rabadan, J., Patel, K., and A. Sajassi, "Updates on EVPN BUM Procedures", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-bum-procedure-updates-14, 18 November 2021, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-bum-procedure-updates-14>.

[I-D.heitz-bess-evpn-option-b] Heitz, J., Sajassi, A., Drake, J., and J. Rabadan, "Multi-homing and E-Tree in EVPN with Inter-AS Option B", Work in Progress, Internet-Draft, draft-heitz-bess-evpn-option-b-01, 13 November 2017, <https://datatracker.ietf.org/doc/html/draft-heitz-bess-evpn-option-b-01>.

[I-D.sr-bess-evpn-vpws-gateway] Rabadan, J., Sathappan, S., Prabhu, V., Lin, W., and P. Brissette, "Ethernet VPN Virtual Private Wire Services Gateway Solution", Work in Progress, Internet-Draft, draft-sr-bess-evpn-vpws-gateway-04, 31 January 2024, <https://datatracker.ietf.org/doc/html/draft-sr-bess-evpn-vpws-gateway-04>.

[I-D.ietf-bess-evpn-ip-aliasing]
Sajassi, A., Rabadan, J., Pasupula, S., Krattiger, L., and J. Drake, "EVPN Support for L3 Fast Convergence and Aliasing/Backup Path", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-ip-aliasing-00, 1 December 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-ip-aliasing-00>.

[I-D.ietf-bess-evpn-mh-split-horizon] Rabadan, J., Nagaraj, K., Lin, W., and A. Sajassi, "EVPN Multi-Homing Extensions for Split Horizon Filtering", Work in Progress, Internet-Draft, draft-ietf-bess-evpn-mh-split-horizon-08, 4 December 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-bess-evpn-mh-split-horizon-08>.

Authors' Addresses

Jorge Rabadan (editor)
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America

Email: jorge.rabadan@nokia.com

Senthil Sathappan
Nokia
520 Almanor Avenue
Sunnyvale, CA 94085
United States of America

Email: senthil.sathappan@nokia.com

Ali Sajassi
Cisco

Email: sajassi@cisco.com

Wen Lin
Juniper

Email: wlin@juniper.net