

INTERNET-DRAFT
Intended Status: Informational
Expires: July 29, 2021

L. Krattiger, Ed.
A. Sajassi, Ed.
S. Thoria
Cisco Systems

J. Rabadan
Nokia

J. Drake
Juniper

January 25, 2021

EVPN Interoperability Modes
draft-krattiger-evpn-modes-interop-03

Abstract

Ethernet VPN (EVPN) provides different functional modes in the area of Service Interface, Integrated Route and Bridge (IRB) and IRB Core connectivity. This document specifies how the different EVPN functional modes and types can interoperate with each other. This document doesn't aim to redefine the existing functional modes but extend them for interoperability.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1	Requirements Language	3
2.	Valid Combinations for Interoperability	3
3.	Service Interface Interoperability	5
3.1	VLAN-Aware Bundle and VLAN-Based	5
3.1.1	VLAN-Aware Bundle Service PE	6
3.1.2	VLAN-Based Service PE	7
3.2	Service Interface Interop Mode of Operation	7
4.	Interoperability for different IRB Types	8
4.1	Asymmetric IRB and Symmetric IRB	8
4.1.1	Asymmetric IRB PE	10
4.1.2	Symmetric IRB PE	10
4.2	IRB Interop Mode of Operation	11
5.	Interoperability for different IRB Core Connectivity Modes . .	12
5.1	Interface-Less and Interface-Ful Unnumbered IRB	12
5.1.1	Interface-Less PE	15
5.1.2	Interface-Ful Unnumbered IRB	15
5.2	Tunnel Encapsulation Consideration	17
6.	Security Considerations	17
7.	IANA Considerations	17
8.	References	18
8.1	Normative References	18
8.2	Informative References	18
9.	Conclusion	19
9.1	Demonstration of Applicability	19
9.1.1	Service Interface Interoperability	19
9.1.2	IRB Types	19
9.1.3	IRB Core Connectivity Types	20

[10](#). Authors' Addresses [20](#)

1. Introduction

Ethernet VPN (EVPN) provides different functional modes in the area of Service Interface, Integrated Route and Bridge (IRB) and IRB connection model. It is understood that the different modes are defined with different use-cases in mind. Even with the specific use-cases and the resulting mode definition, the aim of interoperability is critical.

The following EVPN modes are considered for interoperability. It is limited to most pertinent interop modes as oppose to all permutations. In the future if other modes are identified, it will be addressed in future revisions.

- For Service Interfaces, the VLAN Aware Bundle and VLAN Based types.
- In Integrated Routing and Bridging (IRB) the Asymmetric IRB and Symmetric IRB type.
- Within the IRB connectivity types, interface-less and the interface-ful Unnumbered IRB.

1.1 Requirements Language

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

2. Valid Combinations for Interoperability

The tables below provide an overview of the valid combinations for interoperability described in this Internet-Draft.

For the Service Interface Types as described in [\[RFC7432\] section 6](#) and [\[RFC8365\] section 5.1.2](#). Interoperability considerations are provided for the VLAN-Based Service interface ([\[RFC7432\], section 6.1](#)) and the VLAN-Aware Bundle Service Interface type ([\[RFC7432\] section 6.3](#)). The VLAN Bundle Service Interface ([\[RFC7432\] section 6.2](#)) is not considered at this time.

Table 1 represent the considered Service Interface Types interoperability:

	VLAN-Based	VLAN Bundle	VLAN-Aware Bundle
VLAN-Based	YES	NO	YES
VLAN Bundle	NO	YES	NO
VLAN-Aware Bundle	YES	NO	YES

Table 1

In regards to Integrated Route and Bridge (IRB), two different modes are defined in [\[EVPN-INTERSUBNET\]](#), with [section 5](#) describing Symmetric IRB and [section 6](#) Asymmetric IRB:

The interoperability considerations for Asymmetric IRB and Symmetric IRB mode are represented within this Internet-Draft.

For the IRB Core Connectivity, from all the available modes as described in [\[EVPN-PREFIX\]](#), considered for interoperability is the interface-less mode ([section 4.4.1](#)) in conjunction with only one of the interface-ful modes, namely interface-ful IP-VRF-to-IP-VRF with Unnumbered SBD IRB ([section 4.4.3](#)). With the implementation proximation between the two interface-ful modes, considerations for interoperability between interface-less and interface-ful Numbered are currently not considered. Similarly, the interoperability between the two interface-ful modes is currently not being considered, given the already close relation and to limit permutations. Future revisions of this Internet-Draft might address further variations of interoperability.

Table 2 represent the considered IRB Core Connectivity interoperability.

	Interface-Less	Interface-Ful Numbered IRB	Interface-Ful Unnumbered IRB
Interface-Less	YES	NO	YES
Interface-Ful Numbered IRB	NO	YES	NO
Interface-Ful Unnumbered IRB	YES	NO	YES

Table 2

3. Service Interface Interoperability

3.1. VLAN-Aware Bundle and VLAN-Based

[RFC7432] [section 6](#) describes three different Service Interface Types. The two modes in focus for interoperability are namely the VLAN-Based Service Interface as defined in [\[RFC7432\] section 6.1](#) and the VLAN-Aware Bundle Service Interface as defined in [\[RFC7432\] section 6.3](#). The VLAN Bundle Service Interface is not considered.

The VLAN-Based Service Interface defines an EVPN instance consisting of only a single broadcast domain or "Single Broadcast Domain per EVI" as described in [\[RFC8365\] section 5.1.2](#) Option 1. In this mode, individual BGP Route Distinguisher (RD) and Route Target (RT) are required for each EVI. Each EVI corresponds to a single MAC-VRF identified by the RT, which provides the advantage of an BGP RT constraint mechanisms in order to limit the propagation and import of routes to only the PE that are interested. With VLAN-Based, the MAC-VRF corresponds to only a single bridge table. The VLAN-Based Service Interface uses the EVPN MAC/IP Advertisement route ([\[RFC7432\], section 7.2](#)) with the MUST requirement of the Ethernet Tag ID being set to zero.

Differently, the VLAN-Aware Bundle Service Interface follows a bundling of multiple broadcast domains, with each having its own bridge table, into a single EVI. This refers to the definition of "Multiple Broadcast Domain per EVI" as described in [\[RFC8365\] section 5.1.2](#) Option 2. The advantage of this model is that it doesn't require the provisioning of an RD/RT per broadcast domain, which is a moot point when VLAN-Base uses auto-derivation of RD/RT. With VLAN-Aware Bundle Service, RT Constraint, as defined in [\[RFC4684\]](#), does not help to reduce the dissemination of routes for a BD to the PEs attached to that BD. This is given by the nature of the bundle service where the RT is not sufficient to identify the MAC-VRF and corresponding bridge table. The differences between the two modes of Service Interfaces, namely VLAN-Based and VLAN-Aware Bundle Service Interface, lie in the definition of the Ethernet Tag field in the EVPN routes. While VLAN-Based Service Interface defines the EtherTag as "must be set to zero", the VLAN-Aware Bundle Service interface uses the VID within the EtherTag to identify the bridge table within the MAC-VRF. These two requirements are orthogonal and as a result make the interoperability of the two types mutually exclusive, an interoperability is not achievable (Figure 1).




Figure 1: Interop of different Service Interface Types

As illustrated in Figure 1, the MAC/IP routes exchanged by PE1 and PE2 contain Ethernet Tags 2 and 0 respectively. The receiving PE will not process these routes and will normally discard them (treat-as-withdraw)."

By extending the requirements currently present, an interoperability is achievable. The adjustment would be as follows.

3.1.1. VLAN-Aware Bundle Service PE

In case of VLAN Aware Bundle Service Interface on the receiving PE and with the consideration of VLAN Based Service Interface on the advertising PE:

- MUST Operate in Single Broadcast Domain per EVI.
- Multiple Broadcast Domain per EVI case is not considered.
- MUST allow to send and receive zero EtherTag.
- The import of routes is performed based on the import policy (route-target).

- With single bridge table per MAC-VRF, additional evaluation of the EtherTag field is not required; the bridge table is sufficiently defined by the import route-target.
- No Change to data-plane operation, the MPLS label identifies MAC-VRF + bridge-table, or the VNI identifies the MAC-VRF + the bridge-table.

3.1.2. VLAN-Based Service PE

- Operates in Single Broadcast Domain per EVI.

In case of VLAN Based Service Interface on the receiving PE and with the consideration of VLAN Based Service Interface on the advertising PE:

- Operates in Single Broadcast Domain per EVI.
- MUST allow receiving of non-zero EtherTag.
- No Change in control-plane operation, the EVI import policy (route-target) identifies the broadcast domain (bridge-table) within a MAC-VRF.
- No Change to data-plane operation, the MPLS label identifies MAC-VRF + bridge-table, or the VNI identifies the MAC-VRF + the bridge-table.

While the expansion introduces additional configuration requirement for the VLAN-Aware Bundle Service Interface, it also allows for broader interoperability in the eventuality of Vendor "A" only implemented VLAN-Based while Vendor "B" only implemented VLAN-Aware Bundle Service Interface.

3.2. Service Interface Interop Mode of Operation

When Service Interface interoperability is required, a given PE should follow this section's procedures for all its broadcast domains (BDs) and not just the BDs that need interoperability.

For those BDs where interoperability between VLAN-Aware Bundle and VLAN-Based Service Interface is needed, ignoring the presence of the EVPN routes Ethernet Tag ID on the PEs supporting VLAN-Based mode is not enough. Each PE needs to clearly signal what mode it supports, so that all the PEs attached to the same EVI can understand in what mode the EVI operates.

Consider a scenario where PE1 is attached to the BD range BD1-10 and it operates in VLAN-Aware mode, whereas PE2 is attached to the BD range BD7-20 and operates in VLAN-Based mode. Interoperability is required for the intersecting BDs, I.e., BD7-10.

For PE1, this means BD7-10 need to be separated into a dedicated MAC-VRF each. EVPN routes for each of these four MAC-VRFs MUST be advertised by PE1 with an Ethernet Tag ID of zero. In this way, PE1 indicates the use of VLAN-Based mode for those BDs. On reception, PE1 imports the BD7-10 routes based on the Route Target and ignoring the Ethernet Tag ID, as the Route Target alone is sufficient to identify the correct MAC-VRF and Bridge Table. The remaining BDs on PE1 (range BD1-6) continue operating in VLAN-Aware Bundle mode.

In the same example, other PEs attached to BD1-6 must still process the received Ethernet Tag ID in the EVPN routes from PE1, so that they can identify the correct Bridge Table in a given MAC-VRF.

PE2 operates in VLAN-Based mode for BD7-20, as per [[RFC7432](#)] and [[RFC8365](#)]. PE2's EVPN route advertisements for BD7-20 will include individual Route Targets per BD and an Ethernet Tag ID of zero. On reception, PE2 identifies the MAC-VRF and Bridge Table solely based on the Route Target.

4. Interoperability for different IRB Types

4.1. Asymmetric IRB and Symmetric IRB

The differences in the two inter-subnet forwarding modes, namely Asymmetric IRB and Symmetric IRB, are beyond just the information difference in the control-plane from an EVPN Route Type 2 perspective. The two IRB modes have significant differences in inter-subnet forwarding behavior and as a result different operation during label imposition or encapsulation.

With the Asymmetric IRB mode, the ingress PE performs a "bridge-and-route" operation while the egress PE follows a "bridge-only" approach. Differently, the forwarding behavior in Symmetric IRB mode performs a "bridge-and-route" operation on the ingress PE followed by a "route-and bridge" operation at the egress PE. The significance in difference is not only in the forwarding behavior itself but also around how the respective EVPN attribute are used for driving the inter-subnet operation. More specifically, in the case of inter-subnet forwarding with Asymmetric IRB, MPLS Label1 is used towards the egress PE to specify the MAC-VRF and respective Bridge-Domain for forwarding. In inter-subnet forwarding with Symmetric IRB, MPLS Label2 associated with the IP-VRF is used for the inter-subnet

forwarding operation towards egress PE.

The respective forwarding behaviors are described in [EVPN-INTERSUBNET]. The following steps are required to ensure the interoperability between the Asymmetric and Symmetric IRB modes.




Figure 2: Asymmetric IRB and Symmetric IRB

Figure 2 illustrates the overview of and Asymmetric IRB PE (PE1) and a Symmetric IRB PE (PE2) within a interoperability deployment scenario. Attached to PE1, end-point M1/IP1 is attached to BD0 within MAC-VRF1. Respectively, on PE2 end-point M2/IP2 is connected via attachment circuit to BD2 positioned within MAC-VRF2. IRB0 and IRB2 represent the host-facing IRB interface for inter-subnet communication between the different end-points located in the different IP Subnet. The IRB interfaces for a common MAC-VRF/BD on PE1 and PE2 use the same IP address. With the difference of the IRB modes between PE1 (Asymmetric IRB) and PE2 (Symmetric IRB), there is a difference in the MPLS Label presence as part of the MAC/IP routes exchanged between the PEs. PE1s update contains a single label,

representing MPLS Label1 used for bridging purposes. PE2s advertisement contains two labels, one for bridging and one for routing, as part of the MAC/IP route. While PE1 receives all information necessary from PE2, PE2 is missing information necessary for its routing operation. As a result, Inter-Subnet routing between PE1 and PE2 is not achieved.

By following the current existing forwarding behavior as described in [\[EVPN-INTERSUBNET\]](#), interoperability is theoretically achievable without changes in the control-plane format. Nevertheless, there are steps required that involve predominantly the local behavior of the PE with Symmetric IRB mode.

4.1.1. Asymmetric IRB PE

In case of Asymmetric IRB as the advertising PE and with Symmetric IRB on the receiving PE:

- Asymmetric IRB PE MUST send MAC and IP information with MPLS Label1.

In case of Symmetric IRB as the advertising PE and with Asymmetric IRB on the receiving PE:

- Asymmetric IRB PE MUST be able to ignore MPLS Label2.

4.1.2. Symmetric IRB PE

In case of Symmetric IRB as the advertising PE and with Asymmetric IRB on the receiving PE:

- Symmetric IRB PE has no additional requirements.

In case of Asymmetric IRB as the advertising PE and with Symmetric IRB on the receiving PE:

- Symmetric IRB PE requires to add the host-binding information, MAC and IP, and associates them to the adjacency (ARP/ND) table facing the PE with Asymmetric IRB; this is in addition of adding the MAC address into the MAC-VRF table. Since there is no MPLS Label2 or Route-Target for the IP-VRF, the Host IP is not specifically added to IP-VRF table.

4.2. IRB Interop Mode of Operation

Interoperability between the Asymmetric IRB and Symmetric IRB mode follows specific defined behavior that is predominantly required on the PE that operates in the Symmetric IRB mode. Nevertheless, in support for the interoperability, the PE operating in Asymmetric IRB MUST accommodate the following two minimal requirements (with references to Figure 2): 1) The PE that operates in Asymmetric IRB mode (PE1), MUST send the MAC/IP route including the Host IP address. 2) The PE with Asymmetric IRB (PE1) MUST accept the MAC/IP routes sent from PE2 (Symmetric IRB), while ignoring the additional information of MPLS Label2 and Route-Target of the IP-VRF.

In reference to 1), the PE MUST always send the end-point MAC address, Host IP address and related MPLS Label1 as part of the MAC/IP route towards the PE with Symmetric IRB (PE2). This route will be sent only with MPLS Label1 and the Route-Target of the matching MAC-VRF. In reference to the illustration in Figure 2, PE1 MUST generate and advertise an EVPN MAC/IP route using:

- MAC Length of 48
- MAC Address of M1
- IP Length of 32 / 128
- IP Address of IP1
- Label for MAC-VRF1
- Route-Target of MAC-VRF1
- Next-Hop PE1

For completeness of the requirements and in reference of 2), the MAC/IP route advertised from the PE operating in Symmetric IRB (PE2) is as follow:

- MAC Length of 48
- MAC Address of M2
- IP Length of 32 /128
- IP Address of IP2
- Label for MAC-VRF2, IP-VRF1

- Route-Target of MAC-VRF2, IP-VRF1
- Next-Hop PE2

As defined in 2), the Label and Route-Target information for IP-VRF1 MUST be ignored by PE1 (PE with Asymmetric IRB).

With PE2 operating in Symmetric IRB and with enabled interop mode, the MAC/IP route from PE1 (Asymmetric IRB) is processed in the respective bridging, routing and adjacency table. Based on the Route-Target for MAC-VRF1, the MAC address M1 will be imported into MAC-VRF1 respectively and placed within BD0. In addition, the host-binding information M1/IP1 MUST be installed within PE2s adjacency table. Subsequent, on PE2 the MAC address M1 and the host-binding information (adjacency table entry) of M1/IP1 MUST point towards PE1 as the next-hop. With no presence of the Route-Target for IP-VRF1, the IP address IP1 will not be specifically imported into IP-VRF1 and is not associated with a MPLS Label2. As a result of the interoperability, the additional efficiency provided by Symmetric IRB in regards of preserving adjacency table exhaustion is reduced; this is specifically when communicating with an Asymmetric IRB based egress PE. In contrary, the interop mode allows for communication between the different IRB modes. As a result, in the eventuality that Vendor "A" only provides Asymmetric IRB, while Vendor "B" only has Symmetric IRB available, interoperability for inter-subnet forwarding can be seamlessly achieved. In addition, two further benefits are present by implementing an Asymmetric/Symmetric Co-Existence on the same PE (dual-mode PE).

- A dual-mode PE can seamlessly communicate with PE's that are either in Asymmetric or in Symmetric IRB mode.
- A dual-mode PE can act as Anchor for interconnecting Symmetric IRB and Asymmetric IRB based PE's (deployment restrictions might apply).

5. Interoperability for different IRB Core Connectivity Modes

5.1. Interface-Less and Interface-Ful Unnumbered IRB

The two modes, namely interface-less and interface-ful Unnumbered SBD IRB, are closely related in regards to the information required in the EVPN Route Type 5. While interface-less provides all information for the IP prefix advertisement within the EVPN Route Type 5, in the case of interface-ful Unnumbered SBD IRB, an additional EVPN Route Type 2 is required for the next-hop recursive lookup. From a forwarding behavior, both approaches are similar and follow a symmetric routing approach but are not interoperable. Note as per

[EVPN-PREFIX] the interface-ful Unnumbered SBD IRB is an OPTIONAL mode.




Figure 3: Interoperability of different IRB Core Connectivity Mode (unnumbered)

The illustration in Figure 3 represents the possible deployment scenario between two different Core IRB Connectivity modes. Specifically, PE1 is operating with interface-less Core IRB Mode while PE2 operates with the interface-ful Unnumbered SDB IRB mode; both operate without interoperability capabilities. Attached to PE1 and PE2 respectively, Tenant System 1 (TS1) and Tenant System 2 (TS2) with different IP Subnet are present. TS1 attached on PE1 as well as TS2 attached to PE2 are represented in a common IP-VRF (IP-VRF1), sharing a common Route-Target between the PEs. With the different IRB Core Connectivity modes on PE1 and PE2 respectively, the differences in IP prefix advertisements as described in [EVPN-PREFIX] are present. PE1 advertises only a single EVPN Route Type 5 (IP Prefix Route) for TS1 using the fields following the interface-less mode:

EVPN Route Type 5:

- IP Length of 0 to 32 / 0 to 128

- IP Address of SN1
- Label for IP-VRF1
- GW IP Address set to zero
- Route-Target of IP-VRF1
- Router's MAC Extended Community of PE1
- Next-Hop PE1

Differently, PE2 advertises an EVPN Route Type 2 (MAC/IP Route) next to the EVPN Route Type 5 (IP Prefix Route). The MAC/IP Route supports the requirement for recursive next-hop resolution for the next-hop used in the IP Prefix Route. Below the fields used in the Route Type 5 and respective Route Type 2 according to the interface-ful Unnumbered IRB mode:

EVPN Route Type 5:

- IP Length of 0 to 32 / 0 to 128
- IP Address of SN1
- Label SHOULD be set to 0
- GW IP Address SHOULD be set to "
- Route-Target of IP-VRF1
- Router's MAC Extended Community of PE2
- Next-Hop PE2

EVPN Route Type 2:

- MAC Length of 48
- MAC Address of PE2
- IP Length of 32 / 128
- IP Address of PE2
- Label for IP-VRF1
- Route-Target of IP-VRF1

- Next-Hop PE2

While PE1 is missing the MPLS Label for the IP-VRF from PE2, PE2 is missing the MPLS Label information and the necessary info for the next-hop recursion. As a result, Routing with IP Prefix Advertisement between PE1 and PE2 is not achieved.

By advertising an additional EVPN Route Type 2 from interface-less (PE1) and by advertising the MPLS Label as part of EVPN Route Type 5 from PE2, interoperability is achievable. The specific mode of operation would be as per the following two section and refers to Figure 3 and Figure 4.

5.1.1. Interface-Less PE

In case of interface-less on the advertising PE and with the consideration of interface-ful Unnumbered IRB as the receiving PE:

Shall generate and Advertise EVPN Route Type 2 for every IP-VRF using.

- MAC Length of 48
- MAC Address with "Router MAC"
- IP Length of 32
- IP Address with NVE IP
- Label for IP-VRF
- Route-Target of IP-VRF
- Router-MAC Extended Community

In case of interface-less on the receiving PE and with the consideration of interface-ful Unnumbered IRB as the advertising PE:

- MUST ignore EVPN Route Type 2 with MPLS Label and route-target matching the IP-VRF because there is no MAC-VRF defined matching these information.

5.1.2. Interface-Ful Unnumbered IRB

In case of interface-ful Unnumbered on the advertising PE and with

the consideration of interface-less as the receiving PE:

- Shall advertise MPLS Label for IP-VRF in EVPN Route Type 5 with matching route-target.

In case of interface-ful Unnumbered on the receiving PE and with the consideration of interface-less as the advertising PE:

- No Additions Required.

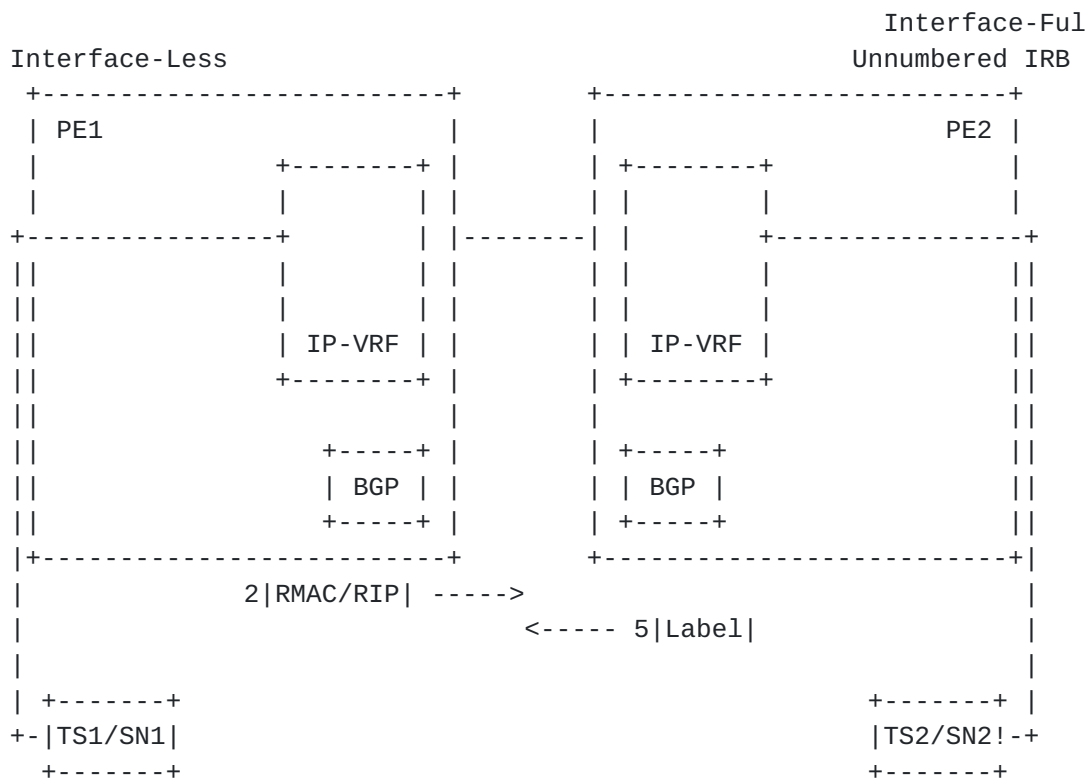


Figure 4: Interop of different IRB Core Connectivity Types (unnumbered)

Illustrated in Figure 4 are the additional requirements for interface-less IRB Core Connectivity mode, specifically the MAC/IP Route (EVPN Route Type 2) necessary for PE2s next-hop recursion. Also, the MPLS Label addition within PE2s IP Prefix route (EVPN Route Type 5) is represented, which is required for interface-ful Unnumbered IRBs advertisement towards an interface-less PE (PE1)

The interop mode introduces additional control-plane advertisements from an Interface-less perspective. This is necessary to allow interface-ful Unnumbered SBD IRB to perform the recursive lookup required. From a EVPN Type 5 perspective between the two types, most

of the fields are already equally defined and populated as per [EVPN-PREFIX]. Exception is the IP-VRF Label, which is required to be added in the interface-ful Unnumbered SBD IRB's EVPN Type 5. In addition, the Interface-less addition allows the Co-Existence of both types on the same PE (dual-mode PE). Such a dual-mode PE can communicate at the same time with PE's that are in Interface-less or in interface-ful Unnumbered SBD IRB mode.

The disadvantage of the additional advertisement has to be put into relation to advantage of successful interoperability where eventually Vendor "A" only implemented interface-less while Vendor "B" only implemented interface-ful Unnumbered SBD IRB.

5.2. Tunnel Encapsulation Consideration

In regards to IRB core connectivity both solutions, namely interface-less and interface-ful, provide a solution for Layer 3 connectivity among the IP-VRFs. Even as the functional result of both modes is the same, there are important considerations in regards to tunnel encapsulations.

[EVPN-IRB] [section 4](#) considers the choice for the NVO tunnel should be dictated by the tunnel capabilities. For example for the IP-VRF-to-IP-VRF model with interface-less, the NVO tunnel for MPLS needs to be IP NVO and for VXLAN needs to be Ethernet NVO.

With the "IP-VRF-to-IP-VRF" model that is used in interface-ful (numbered or unnumbered), [section 4.4.2](#) or 4.4.3 respectively describe the solution to accommodate Ethernet NVO tunnels (VXLAN or GPE, GENEVE, MPLS with MAC payload) only. In the case of interface-ful unnumbered, the Router-MAC Extended Community is always signaled via EVPN update message, which implies the presence of a MAC payload. IP NVO Tunnel are not applicable to these two use-cases/models

Depending on the use of NVO tunnels, interoperability between interface-less and interface-ful unnumbered requires additional changes on the Tunnel Encapsulation mode. This Internet-Draft considers the usage of a compatible NVO Tunnel mode between a PE operating in interface-less and a PE operating interface-ful unnumbered mode.

6. Security Considerations

TBD.

7. IANA Considerations

TBD.

8. References

8.1. Normative References

- [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>>.
- [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>>.
- [EVPN-INTERSUBNET] Sajassi et al., "Integrated Routing and Bridging in EVPN", [draft-ietf-bess-evpn-inter-subnet-forwarding-11](#), work in progress, October, 2020.
- [EVPN-PREFIX] Rabadan et al., "IP Prefix Advertisement in EVPN", [draft-ietf-bess-evpn-prefix-advertisement-11](#), May 2018.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC1776] Crocker, S., "The Address is the Message", [RFC 1776](#), DOI 10.17487/RFC1776, April 1 1995, <<http://www.rfc-editor.org/info/rfc1776>>.
- [TRUTHS] Callon, R., "The Twelve Networking Truths", [RFC 1925](#), DOI 10.17487/RFC1925, April 1 1996, <<http://www.rfc-editor.org/info/rfc1925>>.

8.2. Informative References

- [RFC4684] Marques, P., Bonica, R., Fang, L., Martini, L., Raszuk, R., Patel, K., and J. Guichard, "Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)", [RFC 4684](#), DOI 10.17487/RFC4684, November 2006, <<https://www.rfc-editor.org/info/rfc4684>>.

- [EANTC] EANTC, "Multi-Vendor Interoperability Test", February 2019, <<http://www.eantc.de/fileadmin/eantc/downloads/News/2019/EANTC-MPLSSDNNFV2019-WhitePaper-v1.2.pdf>>.
- [EVILBIT] Bellovin, S., "The Security Flag in the IPv4 Header", [RFC 3514](#), DOI 10.17487/RFC3514, April 1 2003, <<http://www.rfc-editor.org/info/rfc3514>>.
- [RFC5513] Farrel, A., "IANA Considerations for Three Letter Acronyms", [RFC 5513](#), DOI 10.17487/RFC5513, April 1 2009, <<http://www.rfc-editor.org/info/rfc5513>>.
- [RFC5514] Vyncke, E., "IPv6 over Social Networks", [RFC 5514](#), DOI 10.17487/RFC5514, April 1 2009, <<http://www.rfc-editor.org/info/rfc5514>>.

9. Conclusion

With minimal additions, the most common EVPN types for Virtual Identifiers to EVI Mapping, Integrated Routing and Bridging and IP Prefix Advertisement can be made interoperable. The aim for interoperability doesn't remove the requirement for optimized types for different use-cases but allows flexibility and basic interoperability.

9.1. Demonstration of Applicability

Cisco, Juniper and Nokia demonstrated successfully the ability of EVPN interoperability modes during EANTCs yearly "Multi-Vendor Interoperability Test". The Whitepaper can be obtained through EANTC with the latest version being available at [[EANTC](#)].

9.1.1. Service Interface Interoperability

A proof of the benefit with this interoperability mode has already been demonstrated during EVPN Multi-Vendor interoperability testing and also, in production environments. Specifically, Cisco and Nokia's VLAN-Based Service Interface successful proofed interoperability with Junipers VLAN-Aware Bundle Service Interface.

9.1.2. IRB Types

A proof of the benefit with this interoperability mode has already successfully demonstrated during EVPN Multi-Vendor interoperability

testing. Specifically, Cisco operated in a Hybrid IRB (Dual-Mode) mode while other Vendor operated in an Asymmetric IRB mode. Forwarding was achieved through dynamic detection of the alternate Vendor PE's mode and adjustment to Asymmetric IRB for these specific BDs. Communication for all other BDs continued to be Symmetric IRB.

9.1.3. IRB Core Connectivity Types

A proof of an interoperability mode between interface-less and interface-ful Unnumbered SBD IRB has already been demonstrated in production environments and during EVPN Multi-Vendor interoperability testing. Specifically, Cisco's addition for Interface-less is successfully deployed with Nokia's and Nuage's interface-ful Unnumbered SBD IRB at customers

10. Authors' Addresses

Lukas Krattiger
Cisco
USA
EMail: lkrattig@cisco.com

Ali Sajassi
Cisco
USA
EMail: sajassi@cisco.com

Samir Thoria
Cisco
USA
EMail: sthoria@cisco.com

Jorge Rabadan
Nokia
777 E. Middlefield Road
Mountain View, CA 94043 USA
EMail: jorge.rabadan@nokia.com

John E. Drake
Juniper
EMail: jdrake@juniper.net